NAB SHOW 2026

Demo Showcase

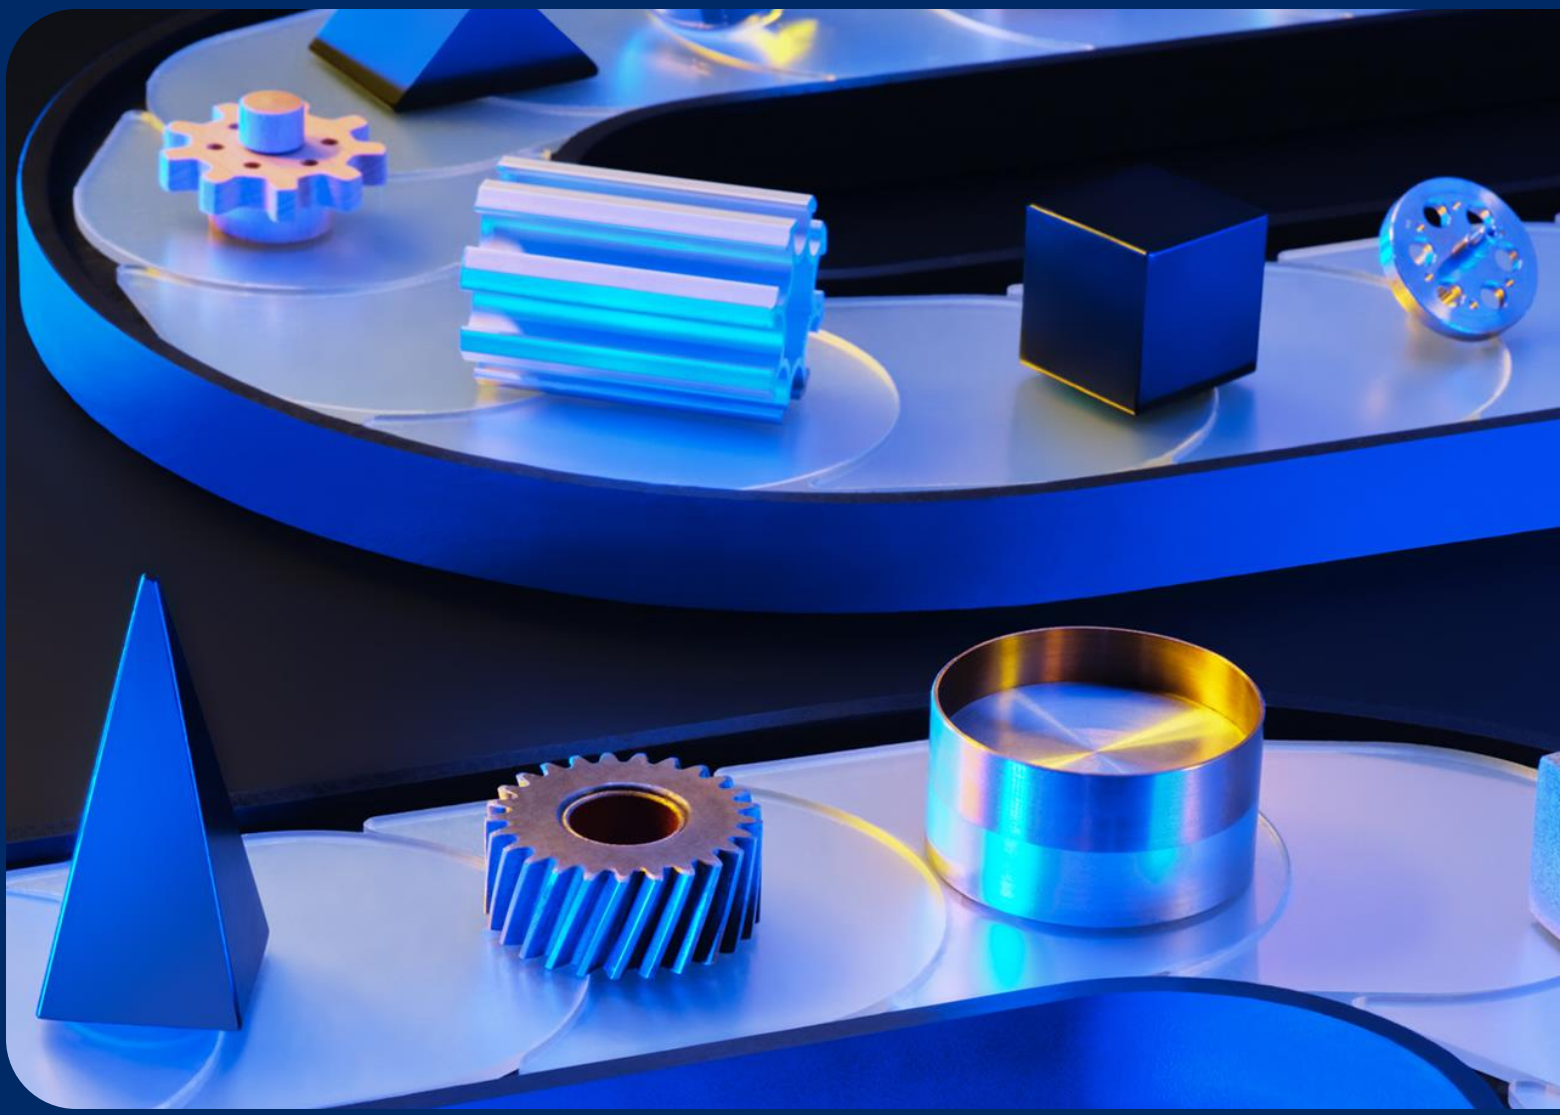

Las Vegas Convention Center
April 19–22, 2026

Demos

Click a solution area to explore demos.

- Content Creation
- Multiplatform Distribution
- Media Lifecycle Management
- Revenue Generation
- Builder Zone

At NAB 2026, AWS shows how cloud infrastructure and agentic AI services remove traditional barriers between creative vision and global reach. 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

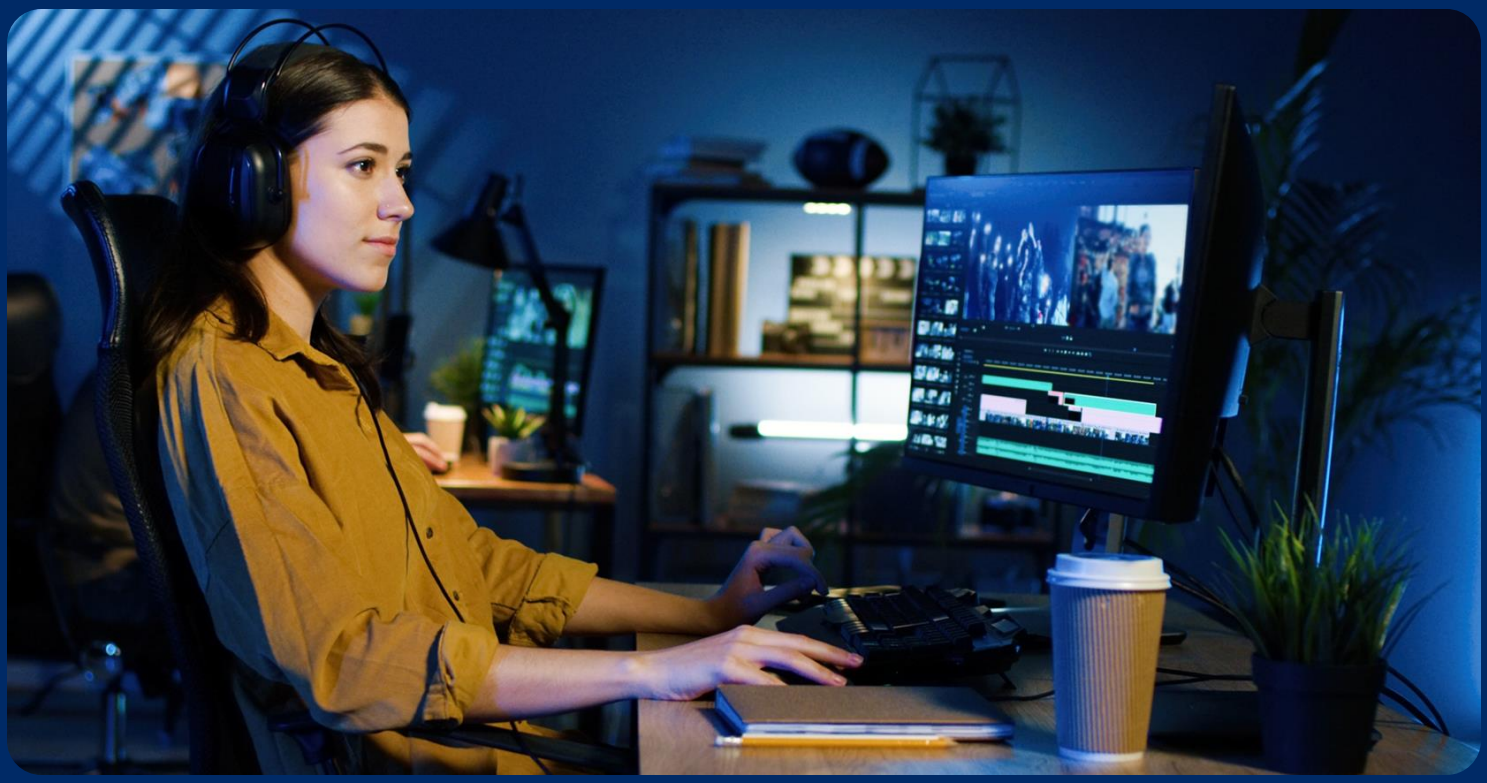

Content Creation

Content Creation demos show how AI-enabled tools and agentic AI workflows empower creators and studios at every production stage. From AI-assisted creative workflow development and cloud-scale rendering with Deadline Cloud to intelligent studio operations and automated podcast production—see how AWS removes bottlenecks, accelerates time-to-market, and delivers flexibility to scale resources on demand across any format or channel.

- Building creative workflow solutions
- Expand and optimize rendering with Deadline Cloud
- Improve studio operational efficiency with AI
- Orchestrating content generation with ComfyUI on AWS
- Automating live audio and video podcast production and distribution

Content Creation

Building creative workflow solutions

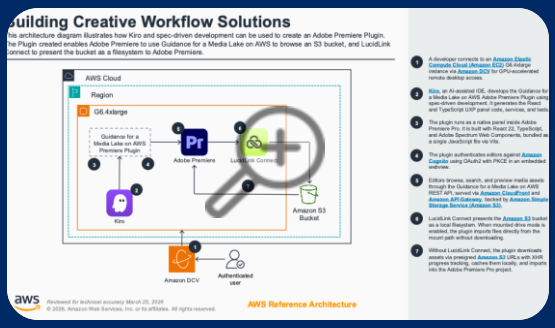

Demo description

Studios struggle to customize creative workflows because building integrations between professional tools often requires specialized development skills, creating bottlenecks and slowing down productivity gains due to limited developer capacity. This demonstration shows how AI-assisted development with Amazon Kiro enables individuals with modest coding backgrounds to rapidly build custom plugins for Adobe Premiere Pro, integrating capabilities like media asset management, ComfyUI workflows and image & video generation directly into artists' existing workspace. Attendees will see live plugin development, including the ease of customization to meet artist needs, demonstrating the productivity gains coding assistants can offer technology teams working with creative artists.

Demo partners

Adobe LucidLink

Related content

- Kiro Getting Started
- LucidLink Connect

Content Creation

Expand and optimize rendering with AWS Deadline Cloud

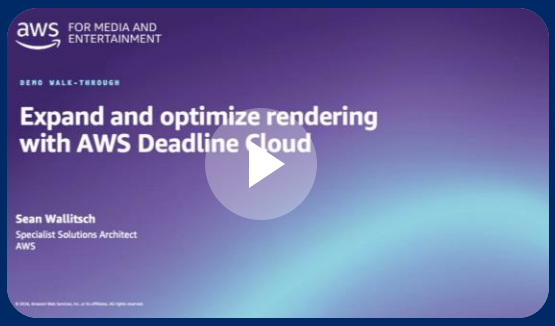

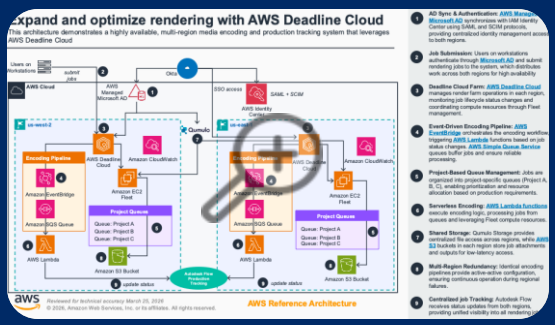

Demo description

Content Creators face long processing times and expensive infrastructure when managing visual effects pipelines. AWS Deadline Cloud helps studios and post-production teams to accelerate their workflows by distributing compute-intensive tasks across cloud resources that scale on demand. In this demonstration we showcase the use of new service capabilities in Deadline Cloud to enable global collaboration and cost optimization strategies that balance rendering speed with budget constraints. Attendees will also experience Deadline Cloud's integration with Nuke CopyCat from Foundry, accelerating common artist workflows such as color corrections, beauty finishes, and deblurring at cloud scale.

Demo partners

FOUNDRY.

Related content

- Deadline Cloud Assistant YouTube
- Deadline Cloud Wait and Save
- Wait and Save Webinar
- Migrate from on-prem to Cloud with Deadline Cloud
- Deadline Cloud Assistant
- Multi-region Deadline Cloud
- Configurable job scheduling
- Host Configuration Scripts
- Cloud Fleet Scaling Configuration
- VPC Lattice with Deadline Cloud

Content Creation

Improve studio operational efficiency with AI

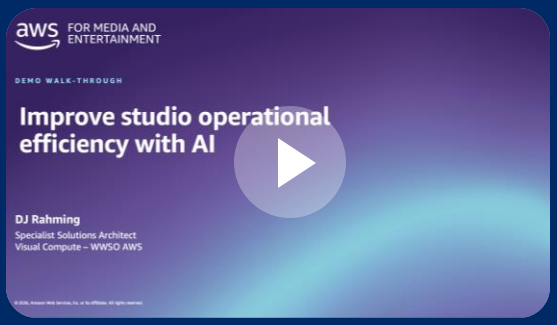

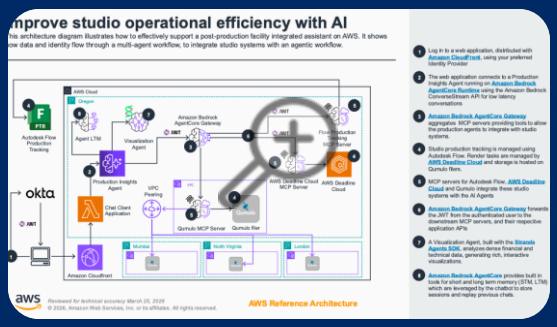

Demo description

Production managers and studio executives often struggle to make time-sensitive decisions about resource allocation, vendor selection, and production logistics due to the complexity of balancing costs, schedules, tax implications, and operational constraints across global facilities. This demonstration shows how AI agents can be applied to analyze cloud production environments, production datasets, compute costs, scripts, and schedules to provide real-time recommendations for complex decisions like relocating VFX work between regions. Attendees will see users ask questions through a natural language interface, receive data-driven recommendations, and then execute approved plans to migrate productions between studios automatically. The result: studios can make faster, more informed decisions that optimize costs and timelines while reducing risk.

Demo partners

AUTODESK Qumulo

Related content

- Autodesk Flow Production Tracking
- Cloud Native Qumulo on AWS
- Amazon Bedrock AgentCore
- AWS Deadline Cloud

Content Creation

Orchestrating content generation with ComfyUI on AWS

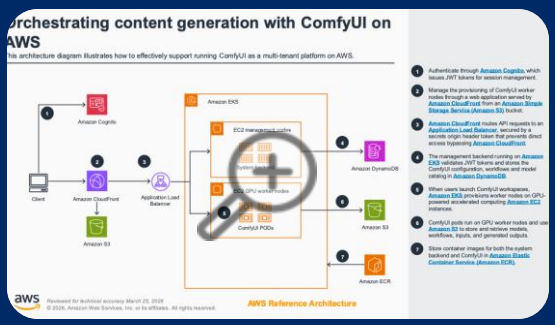

Demo description

Artists working with image, video and audio content generation and transformation tasks need powerful tools like ComfyUI, but struggle with the cost and complexity of maintaining high-performance on-premises infrastructure as well as collaborating with other artists and teams. This demonstration shows how a containerized deployment of ComfyUI enables artists to operate content generation workflows through cloud infrastructure that scales on demand, while leveraging AWS as the collaboration and orchestration hub for content creation. Attendees will see artists working with content generation pipelines—including upscaling, noise reduction, and segmentation—using third-party model providers integrated into the workflow. The result: studios can provide artists with flexible, scalable access to powerful creative AI tools without investing in expensive on-premises hardware or managing complex infrastructure.

Content Creation

Automating live audio and video podcast production and distribution

Demo description

As media organizations increasingly span productions across multiple facilities and regions, collaboration on content scattered across multiple Avid Media Central systems emerges as a challenge. This demonstration shows how Avid Content Core unifies distributed production infrastructure into a single, searchable catalog accessible from anywhere. Attendees will see how Avid workflows can be integrated into a wider production pipeline, powered by Agentic AI and the Content Core API. The result: faster content discovery and streamlined collaboration across global teams working with Avid Media Composer.

Demo partners

ART19 radio.cloud TRITON

Related content

- AWS blog post: Sentiment Analysis with Text and Audio Using AWS Generative AI Services: Approaches, Challenges, and Solutions
- AWS blog post: Building real-time conversational podcasts with Amazon Nova 2 Sonic
- AWS blog post: Building intelligent audio search with Amazon Nova Embeddings: A deep dive into semantic audio understanding

Multiplatform Distribution

Multiplatform Distribution demos showcase how AWS enables broadcast and streaming organizations to deliver the highest quality video across every screen. From agentic news editing and intelligent live event operations to hyper-personalized streaming and low-latency sports delivery—see the unmatched agility, elasticity, scalability, and reliability that reaches new audiences and transforms viewer engagement.

Agentic clipping and assembly for news editing

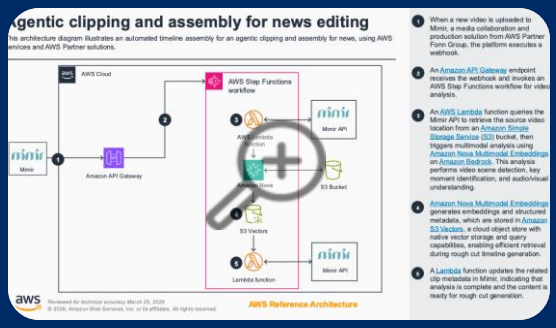

Demo description

News organizations face tight deadlines to identify, curate, and assemble relevant footage for breaking stories. AWS addresses this challenge through Amazon Bedrock and Amazon Nova foundation models, enabling intelligent content analysis and automated timeline assembly. This demonstration showcases how agentic AI analyzes news scripts and existing footage to generate recommended rough-cut timelines with contextually relevant clips. By combining Amazon Nova's multimodal understanding with automated editing workflows and AWS Partner Solutions from Fonn Group's Mimir and Saga, news teams accelerate story production while maintaining editorial control and can focus on storytelling rather than asset management.

Demo partners

mimir saga

Related content

- AWS for M&E Blog: Edit in the cloud: Nine Network modernizes edit suites with AWS

Agentic intelligent operations for live events

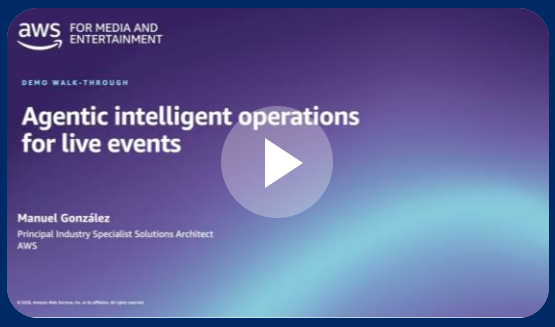

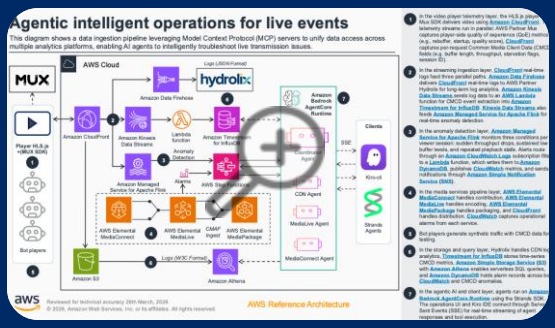

Demo description

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 the AWS Strands SDK, this demonstration showcases four specialized AI agents designed to identify root causes of technical issues end-to-end across broadcast and OTT workflows. By integrating observability data from AWS Media Services, CMCD video player data and streaming analytics from AWS Partners Mux and Hydrolix, 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

hydrolix MUX

Related content

- 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

AI-driven archive monetization and localization

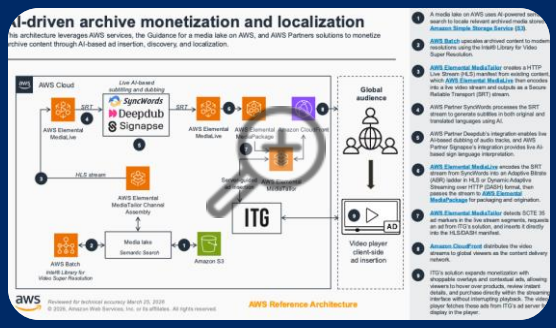

Demo description

Media companies hold vast libraries of valuable content, but struggle to monetize it effectively. Organizations often lack clear understanding of their archive content due to outdated formats, missing or incomplete metadata, inferior content discovery abilities, and highly manual and costly legacy workflows. This demonstration shows how AWS enables intelligent content discovery, automated metadata enrichment, and new revenue streams from archive libraries. Using generative AI, organizations can enhance archive video resolution for modern screens, enable advertising-supported streaming with intelligent ad insertion, and automate localization through subtitles and dubbing for global markets. Attendees will see how to transform dormant archives into revenue-generating assets that reach new geographies and broader audiences.

Demo partners

deepdub itg Signapse SyncWords

Related content

- AWS Blog: Configure live captions with SyncWords and AWS Elemental Media Services
- AWS Blog: 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
- AWS Blog: Delivering low-latency captions and voice translation for live sports, news, and OTT platforms with Syncords and AWS
- AWS Blog: 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
- AWS Blog: Live streaming localization and accessibility using AWS Media Services

AI-driven contextual advertising for live news broadcasts

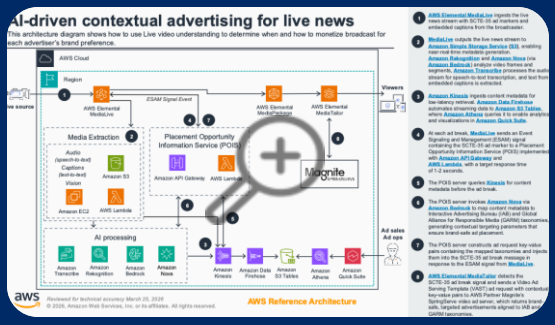

Demo description

Live news broadcasters struggle to maximize ad revenue while maintaining brand safety due to the fast paced and dynamic nature of a live news broadcast. This demonstration shows how Amazon Bedrock and AWS Elemental Media Services can be combined to analyze live video streams in near real-time, detecting people, places, objects, and key moments to generate time-addressable metadata aligned with IAB and GARM brand-safety frameworks. The result: news broadcasters can offer premium, contextually relevant ad inventory at scale while ensuring brand safety compliance.

Demo partners

Magnite

Related content

- IABM Content Taxonomy

AWS Elemental MediaConnect Router: Simplifying live contribution workflows

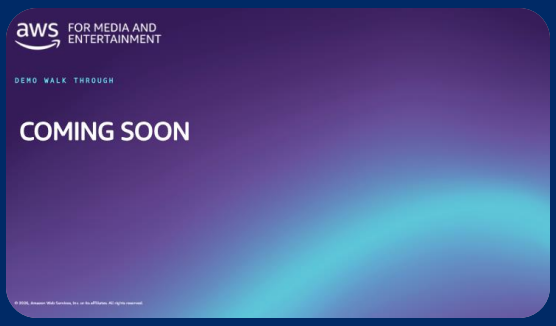

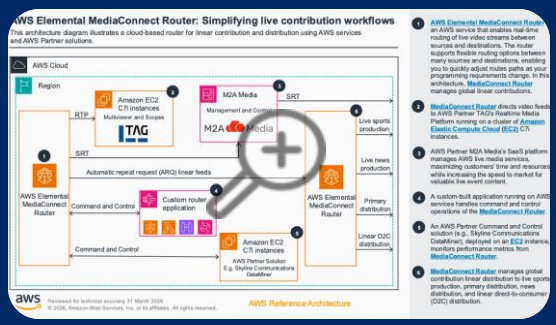

Demo description

As broadcasters modernize live contribution workflows, the cloud offers a powerful opportunity to simplify the routing, switching, and monitoring of multiple live feeds across globally distributed sources. This demonstration shows how AWS Elemental MediaConnect Router enables centralized management of live and linear feeds from contributors worldwide. Attendees will see real-time feed routing and switching through a physical remote control panel (RCP), a custom-built operator interface, and the native AWS Elemental MediaConnect Router console. Integrated monitoring is provided by an AWS Partner Solution from TAG Video Systems, providing continuous visibility into feed health and automated alerting. The result: broadcasters can bring traditional switching capabilities to the cloud while reducing infrastructure complexity and gaining global operational flexibility.

Demo partners

M2A Media

TAG VIDEO SYSTEMS

Related content

- AWS Elemental MediaConnect Router official page
- AWS Announces Elemental MediaConnect Router
- Using the MediaConnect router

Multiplatform Distribution

Dynamic routing and AI assistants for live cloud production

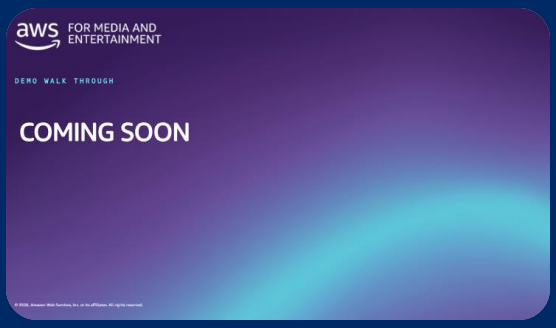

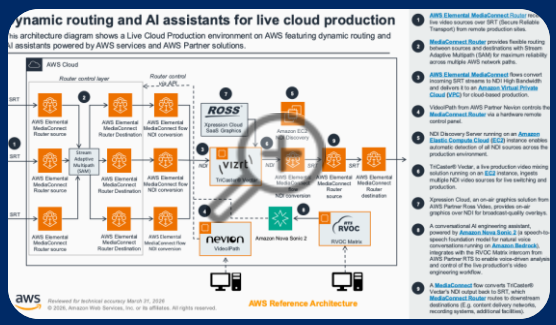

Demo description

Cloud-based live production gives broadcasters the ability to do more with less operational overhead. This demonstration showcases how broadcasters can leverage AWS to enable intelligent automation and streamlined operations at every stage of production. AWS Elemental MediaConnect Router enables dynamic source routing within a scalable, cloud-native infrastructure—including support for global sources—while a dedicated MediaConnect flow handles video-over-IP NDI conversion. This demonstration also showcases a conversational AI engineering assistant, integrated directly into a broadcast intercom system and powered by Amazon Nova Sonic 2, enabling real-time, voice-driven interaction with the live production workflow and environment. Attendees will see how combining AWS Media Services with voice AI accelerates decision-making on the production floor and opens new possibilities for leaner, more agile live broadcast operations.

Demo partners

nevion ROSS RTS RVOC vizrt

Related content

- AWS Elemental MediaConnect Router
- Contribute content to AWS using NDI with AWS Elemental MediaConnect
- Introducing Amazon Nova 2 Sonic: Our new speech-to-speech model for conversational AI
- Guidance for Programmatic Deployment of NDI Discovery Servers for Broadcast Workflows on AWS

Fast-turnaround workflows using TAMS

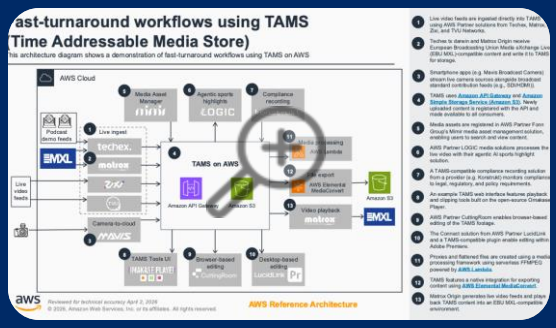

Demo description

Fast-turnaround production teams in news, sports, and live entertainment face pressure to deliver content from camera to audience at unprecedented speeds, but traditional workflows rely on costly or proprietary systems. The open-source Time-addressable Media Store (TAMS) API on AWS enables content-centric, cloud-native workflows from live video and direct camera ingest through editing, AI-powered automated highlight generation, MAM integration, and export for VOD and social media. The demonstration showcases MXL-to-TAMS interoperability and an end-to-end cloud workflow where media is stored once and immediately accessible across the entire production pipeline.

Demo partners

Adobe · CuttingRoom · LOGIC · LucidLink · matrox video · mimir · techex. · TVU networks · ZIXI

Related content

- TAMS Website
- TAMS API Specification (BBC Github)
- AWS Open Source TAMS API Implementation
- AWS Open Source TAMS Tools
- Guidance for Cloud-Native Fast-Turnaround Media Workflows on AWS
- Reuters and AWS demonstrate next generation News distribution at IBC 2025

Hyper-personalized streaming experiences

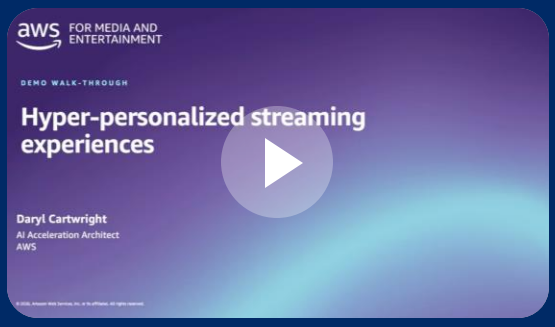

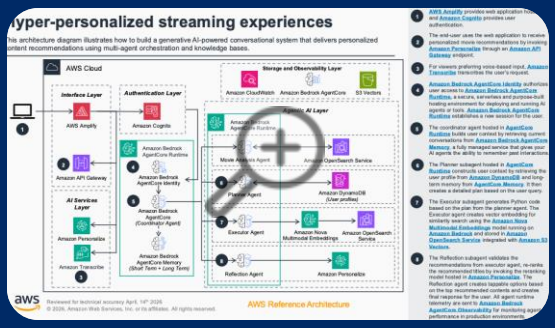

Demo description

The explosion of content across countless channels and platforms has created a paradox: viewers have unlimited choice but struggle to find what they truly want to watch. AWS addresses the dual challenge of audience fragmentation and personalized engagement through Amazon Bedrock and agentic AI, enabling conversational, intent-aware systems that understand natural language and viewer preferences at scale. This demonstration showcases how a combination of traditional machine learning, multi-agent orchestration, and multi-source knowledge bases can deliver a real-time conversational discovery system. By combining intelligent content recommendation, semantic understanding, and real-time content matching, streaming services can deliver personalized experiences that reduce decision paralysis and increase viewer engagement.

Demo partners

IMDb

Related content

- AWS GitHub Sample: Speech to Speech Movie Recommendation Solution
- AWS Blog: 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?

Live video intelligence with TAMS and agentic AI

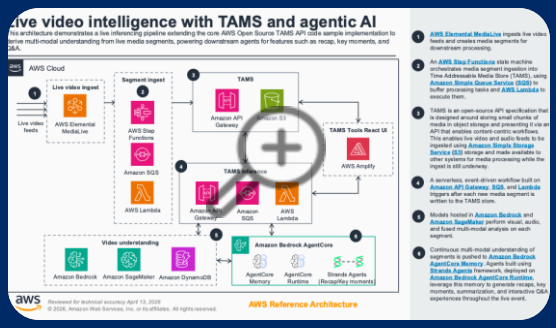

Demo description

Live event coverage generates high-value moments that demand immediate distribution, yet traditional workflows require hours of manual review and editing before media reaches audiences. This demonstration shows how AWS enables real-time multimodal analysis of live video streams to streamline fast-turnaround clip creation workflows. AI Agents continuously analyze live feeds to generate rolling summaries, extract broadcast-quality clips, and assemble highlight reels of key moments as events unfold. By combining Time Addressable Media Store (TAMS) and multiple AI Agents working in parallel, broadcasters can compress traditional multi-hour production workflows into minutes to deliver impactful moments to audiences as they happen.

Related content

- AWS Workshop: Video Understanding with Generative AI on AWS
- TAMS Website
- TAMS API Specification (BBC Github)
- AWS GitHub: Open Source TAMS API Implementation

Multiplatform Distribution

Low-latency streaming and personalized audio for live sports

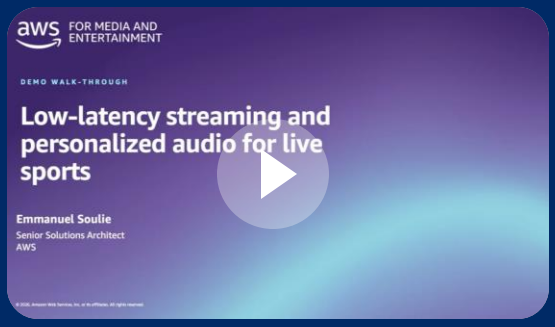

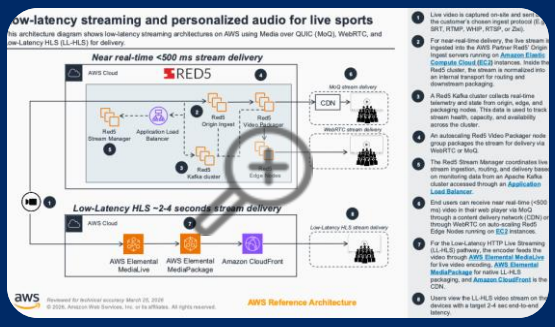

Demo description

Live sports fans at home have always experienced a delay compared to being in the stadium, frustrating fans who want real-time experiences as close to the action as possible. This demonstration shows how protocol advances and cloud-based audio innovation are closing the gap between in person and at home experiences. The demonstration showcases how Low-Latency HLS, WebRTC, and Media over QUIC (MoQ) are able to reduce latency to below 2 seconds. Simultaneously, object-based audio technology on AWS enables viewers to personalize their listening experience -- bringing the crowd, commentary, and atmosphere to life in an entirely new way. Attendees will see how these technologies combine to deliver premium live sports experiences that rival in-stadium viewing while adding personalization that was previously impossible in traditional broadcasts.

Demo partners

jünger MAINCONCEPT RED5 techex.

Related content

- AWS Blog: How to configure a low-latency HLS workflow using AWS Media Services
- IETF Media Over QUIC (moq) working group
  MPEG-H Audio
- Red5: MOQ Streaming overview
- Unlocking next generation audio production with MPEG-H on AWS

MXL: Multi-vendor live media interoperation

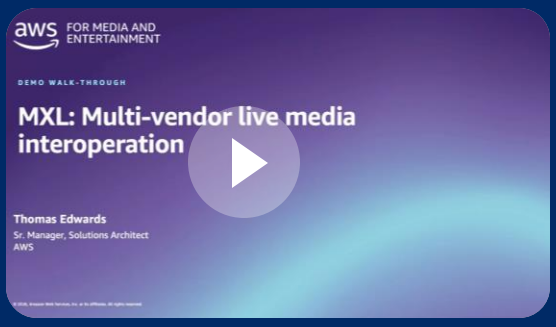

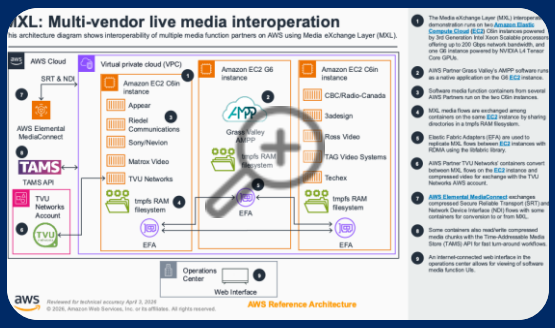

Demo description

Cloud-based live production often forces broadcasters to choose between vendor lock-in or complex integration challenges. The Media eXchange Layer (MXL) is an open-source SDK enabling interoperable ultra-low latency exchange of live uncompressed video, audio, and metadata between different software media functions. This demo showcases MXL v1.0 exchanging live media between multiple AWS Partners hosted on three Amazon EC2 g6 instances using both shared memory within an instance and Amazon Elastic Fabric Adapter (EFA) RDMA transport between instances. Attendees will see how MXL delivers broadcast-quality performance while maintaining vendor choice and workflow flexibility for modern cloud-based live production.

Demo partners

Grass Valley · matrox video · nevion · ROSS · TAG Video Systems · techex. · TVU networks

Related content

- European Broadcasting Union: Media eXchange Layer
- GitHub Repo: MXL : Media eXchange Layer
- VIDTRANS 2026: Media eXchange Layer MXL Technical Update - Thomas Edwards, AWS
- EBU DMF Media eXchange Layer (MXL): Streamlining Multi-Vendor Live Video
- Grass Valley: JT-DMF and MXL: Shared Foundations for Software-Defined Production at Scale
- How Media Exchange Layer (MXL) Could Reshape Software-Based Broadcast Workflows
- TV Tech: MXL Moves into Stable Production Release
- NewscastStudio: Understanding the Media Exchange Layer: A shift in how broadcast software connects

Near-live highlights, every format, every screen with AWS Elemental Inference

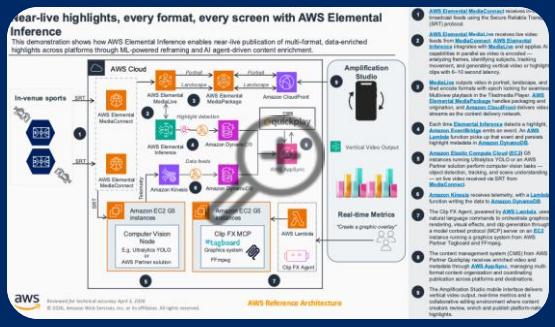

Demo description

Audiences increasingly consume content on platforms broadcasters don't control, from TikTok to Instagram Reels to YouTube Shorts, and they expect to see content optimized for each platform. Manual reformatting is too slow, and post-production AI tools add complexity and delay to publishing workflows. This demonstration shows how AWS Elemental Inference, integrated with AWS Elemental MediaLive, enables and optimizes near-live publication across formats and platforms. First, live action becomes publish-ready vertical video within seconds through ML-powered detection, tracking, and intelligent reframing that runs in parallel with encoding. Second, an AI agent, controlled with natural language within a broadcast-grade editing interface, enriches the content with performance insights, contextual data overlays, and visual effects. Together with AWS Partner Solutions, these stages transform a live broadcast into multi-format, data-driven content that deepens fan engagement without delay.

Demo partners

HAWK-EYE INNOVATIONS · quickplay · #tagboard · tiled media

Related content

- AWS Elemental Inference
- New AI Tool From Amazon's AWS Speeds Up Vertical Video Clip Process; Fox, NBCUniversal Among Initial Customers
- Transform live video for mobile audiences with AWS Elemental Inference

Reinventing linear distribution on AWS

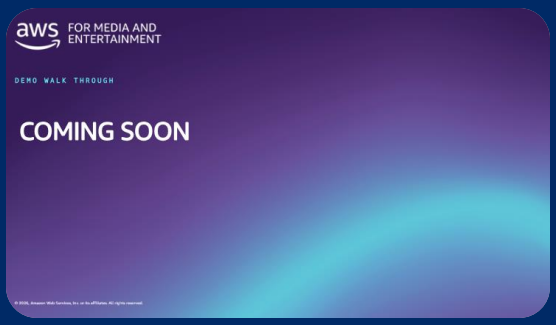

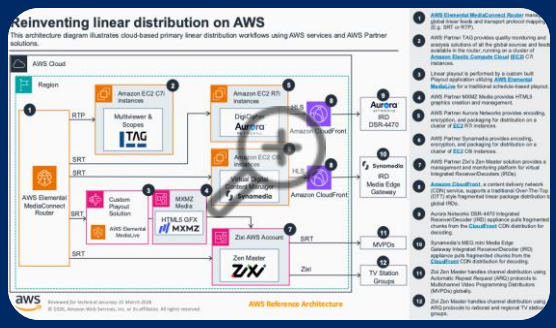

Demo description

The C-band spectrum reclamation is prompting US broadcasters to reassess how they deliver linear channels, and for many, it presents an opportunity to rethink distribution strategy from the ground up. By navigating this change, broadcasters can build a more flexible, cost-efficient approach that utilizes the AWS global infrastructure, including services like AWS Elemental MediaConnect and Amazon CloudFront, enabling multichannel distribution and moving away from satellite dependency and toward cloud-native delivery. This demonstration also showcases a custom-built playout solution driven by Broadcast Exchange Format (BXF), using AWS Elemental MediaLive as the underlying engine. Together, these components illustrate how reliable linear distribution can be maintained and improved through the transition to AWS. With this solution, broadcasters can achieve reliable, cost-effective, broadcast-grade distribution on AWS. Attendees will learn how to architect scalable distribution pipelines suited to any broadcaster looking to augment or replace satellite infrastructure.

Demo partners

Synamedia TAG VIDEO SYSTEMS ZIXI

Related content

- Global Cloud-Based Live Video Delivery with AWS and Zixi
- AWS Elemental MediaConnect Flows Managed & Monitored in ZEN Master
- AWS Blog: How Synamedia delivers an optimized on-demand viewing experience with Cloud DVR on AWS
- 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

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.

AI-assisted promo assembly with creative control

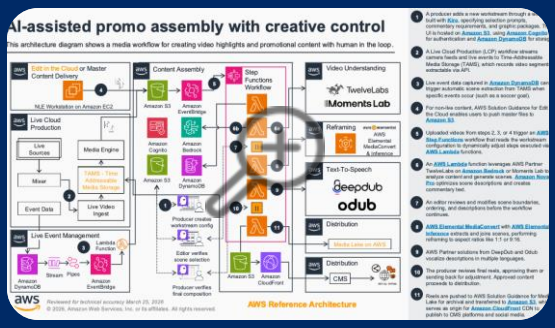

Demo description

Media companies and rights holders struggle to scale promotional content creation—highlights, trailers, and reels—across multiple platforms, formats, and languages while maintaining creative quality. Traditional workflows require significant manual effort to select scenes, adapt formats, and coordinate localization, making it difficult to scale without compromising editorial standards.

This demonstration showcases a content composition platform that amplifies human creativity rather than replacing it. Powered by Amazon Bedrock with Amazon Nova foundation models, Amazon Elemental Inference, and AWS Partner Solutions from MomentsLab, TwelveLabs, and Deepdub, the system uses natural language prompts and AI agents to orchestrate scene selection, localized narration, and context-aware reframing. Human oversight is maintained at three decision points: workstream configuration, scene review, and final output approval.

The result: faster campaign creation, scalable multi-platform delivery, and creative teams focused on direction rather than repetitive production tasks.

Demo partners

deepdub · MomentsLab · odub · TwelveLabs

Related content

- TwelveLabs video understanding models are now available in Amazon Bedrock
- Automated highlights, reels creation and discovery at NAB 2025

Media Lifecycle Management

AI-driven content understanding and metadata creation

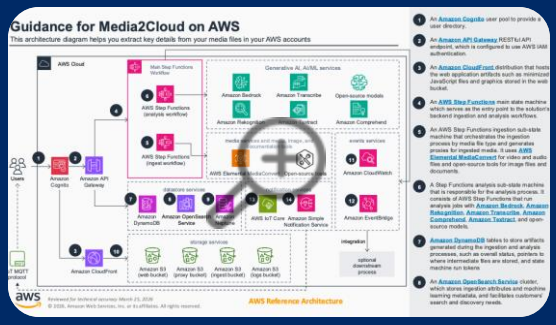

Demo description

Media companies struggle to unlock value from vast video archives, where valuable content remains undiscoverable without costly manual tagging. Finding specific scenes, identifying talent appearances, detecting ad breaks, or locating reusable footage requires hours of review, delaying production and limiting monetization opportunities. Media2Cloud V5 transforms unstructured video into searchable assets, demonstrating cost-effective detection of segments, credits, and reusable clips without compromising accuracy. Guidance for a Media Lake on AWS then provides a centralized, searchable repository that organizes this enriched metadata alongside the source content for discovery and reuse at scale. The media lake pipelines then processes video files to generate frame-level insights, transforming hours of manual work into seconds of automated analysis. The result enables teams to instantly search through years of archives, accelerating content reuse and unlocking new revenue streams from previously inaccessible material.

Related content

- Media2Cloud on AWS Guidance: Scene and ad-break detection and contextual understanding for advertising using generative AI
- Guidance for Media2Cloud on AWS
- Media2Cloud GitHub

AI-powered advertisement compliance for media content

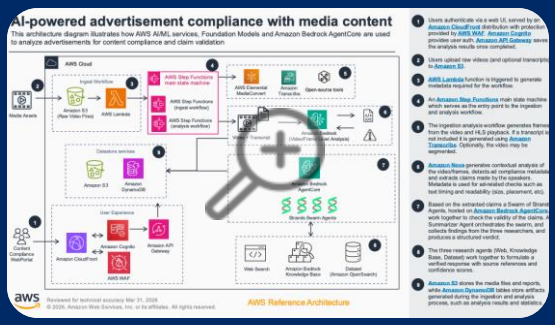

Demo description

Media organizations face mounting pressure to validate advertisements across diverse campaigns—from political campaign ads with funding disclosures and claim verification to pharmaceutical disclaimers, automotive safety assertions, and consumer products. This AI-powered advertisement compliance demonstration showcases how video understanding and agentic AI automate the analysis and verification of commercial advertisements.

Powered by Amazon Bedrock and Amazon Bedrock Agentcore, a swarm of AI agents work together to perform accelerated automated claim validation, fact-checking, and compliance verification against regional regulations and organizational policies. This solution reduces manual review bottlenecks, accelerates ad approval workflows, and ensures consistent adherence to broadcast standards, enabling media organizations to efficiently manage high-volume advertising while maintaining editorial integrity and regulatory

AI-powered localization at scale: Breaking through resource and time barriers

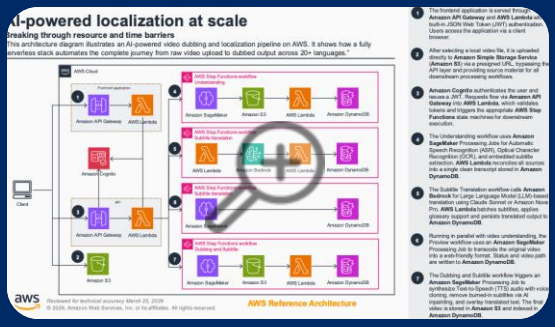

Demo description

Scaling global content delivery without compromising quality and authenticity continues to be a challenge for media organizations. Localization teams struggle with scarce voice talent, fragmented workflows across continents, and time-to-market pressure. This demonstration shows a complete localization pipeline, including the transformation of a production into multiple language versions, preserving tone and emotion and the use of AI-assisted dubbing platforms, supporting synthetic voices and voice cloning, to create local versions of the production with respect for the original performances. The result: content teams can create localization that balances automation speed with emotional authenticity — from premium productions requiring human nuance to high-volume catalogs where intelligent automation delivers quality at scale.

Demo partners

deepdub odub TwelveLabs

Related content

- DeepDub's AI Localization tools unlock business opportunities for Paramount

Content compliance with agents

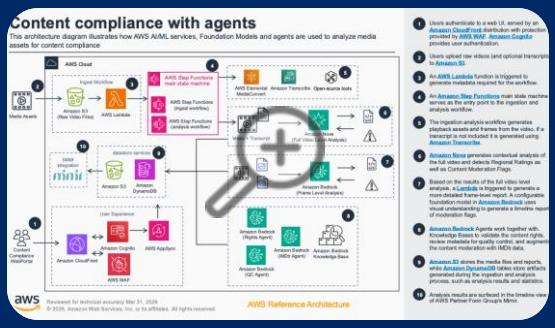

Demo description

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

mimir

Related content

- Streamlining content compliance: Automating media analysis with Amazon Nova

Media Lifecycle Management

Multimodal media discovery and metadata management

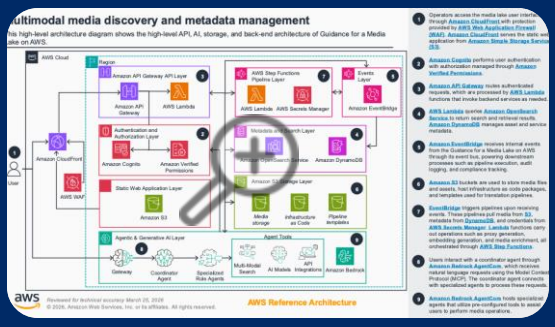

Demo description

Media companies struggle to find and reuse content buried in vast video archives spread across multiple storage locations. This demonstration builds on the Guidance for a Media Lake on AWS to showcase a unified, intelligently searchable catalog of media assets that brings together multiple storage locations behind a single logical media lake. It showcases the use of Amazon Nova MME and AWS Partner TwelveLabs' foundation models to deliver advanced multimodal video understanding, enabling users to perform rich semantic and natural language searches over their media. Integrations with upstream metadata providers, such as the Guidance for Media2Cloud on AWS, enrich foundational technical and descriptive metadata, while the core focus remains on intelligent content understanding that captures what is happening within the media itself. By combining workflow management, metadata enrichment, and multimodal understanding in one system, this solution transforms media workflows—allowing content creators and media professionals to discover, analyze, and extract actionable insights from their libraries, rather than merely searching for files.

Demo partners

TwelveLabs

Related content

- 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
- Guidance for a Media Lake on AWS

Transform media contracts into actionable metadata with vision AI

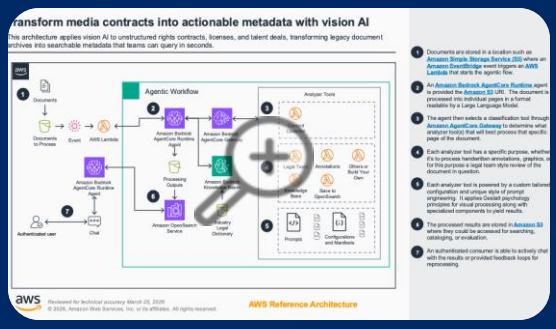

Demo description

Media companies operate on a web of agreements — rights contracts, distribution licenses, talent deals, and guild obligations — yet the terms governing what you own, what you owe, and what you can monetize are buried in PDFs, email threads, and legacy archives accumulated over years. When a business decision depends on understanding contract terms — such as whether a delivery method is specified or when a license window expires — teams resort to manual document review, slowing operations and introducing risk. This demonstration shows how multi-modal AI models extract and structure contract metadata at scale, transforming unstructured documents into searchable, actionable data. Attendees will see the system process diverse contract formats, identifying territory grants, compensation structures, expiration dates, and clause relationships. The result enables teams to search years of contract history in seconds, turning inaccessible archives into strategic assets for rights management and deal optimization.

Related content

- Broad Agentic Document Generative Extraction & Recognition System (BADGERS) AWS Sample Project

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 ad creation servers and agentic advertising workflows to prebid advertising for live sports and automatic localization for global markets—see how agentic AI-powered automation maximizes revenue across advertising, subscriptions, and content licensing.

- Ad creation server
- Agents for advertising
- Maximizing live sports revenue through prebid and contextual non-intrusive ads
- Unlocking new markets with automatic localization for live sports

Revenue Generation

Ad creation server

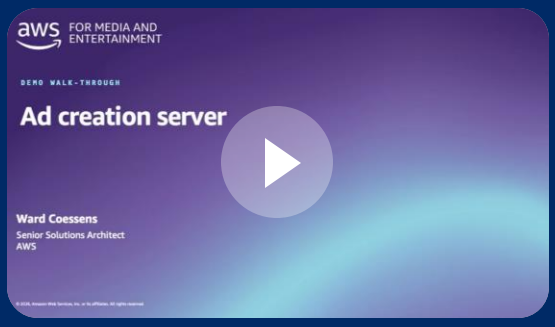

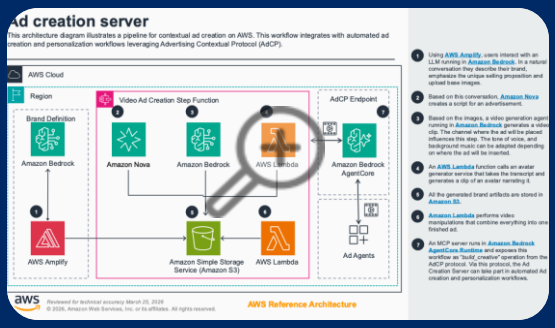

Demo description

Media publishers and advertising platforms need accessible ways to help small and medium businesses (SMBs) create television advertisements without requiring production expertise or large budgets. This demonstration shows how business owners can generate professional ads through natural language chat by simply describing their brand and goals while publishers provide ad placement context, allowing ads to be tailored to specific slots at minimal cost. Attendees will see how an agentic AI workflow, powered by Amazon Bedrock, Amazon Nova models, and leveraging AdCP open agentic advertising standards, guides users from concept to finished ad in minutes. The result: media companies can tap into the (SMB) market, creating new revenue streams from customers that previously focused on their ad spend elsewhere.

Related content

- Leveraging Amazon Nova Reel for Ad creatives
- Supercharging Ad Creative with Amazon Bedrock and Amazon Nova

Revenue Generation

Agents for advertising

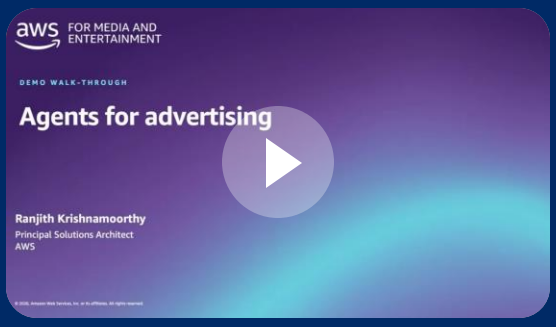

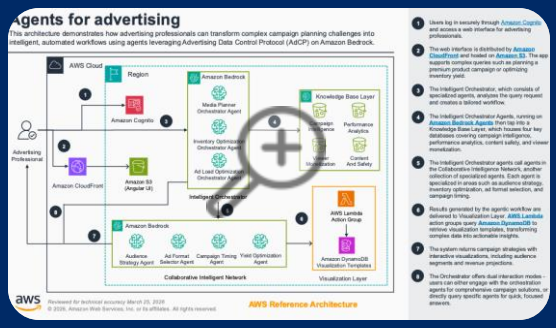

Demo description

Advertising professionals struggle with time-intensive campaign planning that lacks strategic depth and often introduces costly errors. Manual optimization processes are prohibitively complex to manage at scale. This demonstration shows how agentic AI can transform the entire advertising workflow—from media plan generation and audience targeting, to real-time bid optimization and ad monetization. Built on Amazon Bedrock AgentCore and the Strands SDK, and leveraging the AdCP open agentic advertising standards, attendees will see multi-agent collaboration across the advertising value chain. The result: faster campaign planning, improved ad yield, and reduced operational complexity through intelligent automation.

Related content

- Solution Guidance

Revenue Generation

Maximizing live sports revenue through prebid and contextual non-intrusive ads

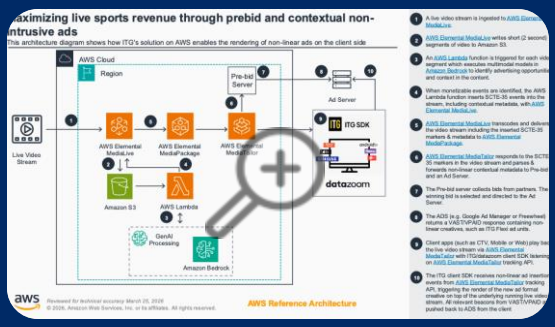

Demo description

Publishers and broadcasters struggle to maximize ad revenue while keeping audiences engaged with innovative, non-intrusive ad experiences. This demonstration shows how to optimize real-time bidding auctions at scale through an integrated solution combining Prebid Server with AWS RTB Fabric, while delivering interactive, shoppable ads powered by AWS Elemental MediaTailor integrated with AWS Partner InTheGame. Attendees will see parallel bid requests handled at scale, with full transparency across multiple demand sources, alongside contextual product placements that enable instant purchases through QR codes or companion apps. The result: increased ad yield through competitive bidding, new revenue streams beyond traditional spots, and enhanced viewer experience with relevant advertising that drives measurable e-commerce outcomes—all while maintaining direct control over demand connections and bidding data.

Demo partners

itg

Related content

- Automate video insights for contextual advertising using Amazon Bedrock Data Automation
- Contextualized viewer engagement and monetization for live OTT events
- Guidance for Deploying a Prebid Server on AWS

Revenue Generation

Unlocking new markets with automatic localization for live sports

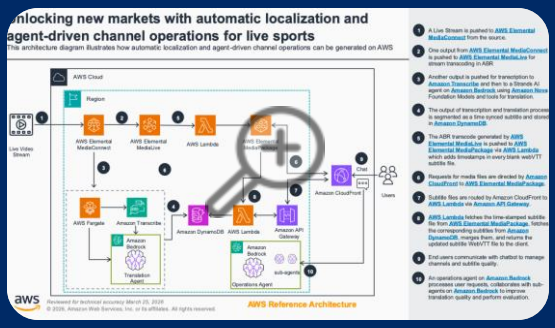

Demo description

Live sports have the power to unite global audiences—but reaching new markets requires more than just delivering the game. Today's fans expect real-time, culturally resonant, and localized experiences in their preferred language. This demonstration shows how automatic speech-to-text conversion and an agentic AI translation workflow, powered by Amazon Bedrock and AWS Elemental Media Services, generate accurate, low-latency subtitles in multiple languages simultaneously during live sporting events, preserving sports-specific terminology while ensuring cultural appropriateness. Attendees will see a live sports stream being transcribed and translated in real-time, with subtitles appearing in multiple languages within seconds of the original commentary. The result: broadcasters can expand into new markets without the cost and complexity of manual localization teams.

Related content

- AWS Blog: Scale global live reach with AWS powered real time webvtt multilingual subtitling

Builder Zone

Builder Zone demos showcase cutting-edge AWS innovations that push the boundaries of media technology. From AWS Elemental Inference for AI-powered vertical video and automated quality monitoring to cloud-native video routing at scale and secure streaming delivery—explore next-generation solutions that remove technical barriers between creative vision and global audience reach.

- AWS Elemental Inference — AI-powered vertical video in real-time
- Automated quality monitoring and failover for live streaming
- Cloud-native video routing at scale
- Secure streaming delivery using Common Access Tokens and CBOR Web Tokens

AWS Elemental Inference – AI-powered vertical video in real-time

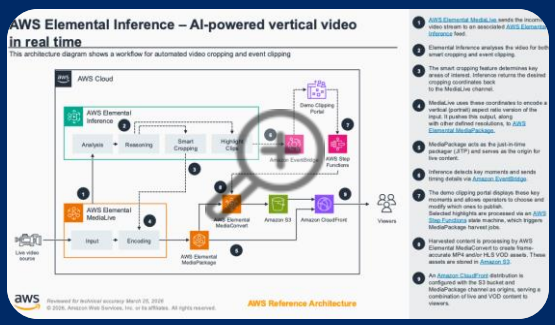

Demo description

Broadcasters and content creators streaming events to vertical-first platforms such as TikTok and Instagram are challenged by the need to intelligently frame vertical video from live sources while also tagging and extracting highlight-worthy moments to deepen fan engagement. This demonstration shows how AWS Elemental Inference performs AI-powered live video analysis in parallel with existing workflows. Attendees will see the generation of broadcast-quality vertical video from live sources, in real-time, along with highlight clipping for posting to social media and website platforms. The result: simultaneous multi-format distribution and fast turnaround social media clips, all from a single live source without workflow interruption.

Related content

- AWS Elemental Inference Documentation
- AWS Elemental Inference FAQ
- Smart cropping video using Elemental Inference
- Clipping video using AWS Elemental Inference

Builder Zone

Automated quality monitoring and failover for live streaming

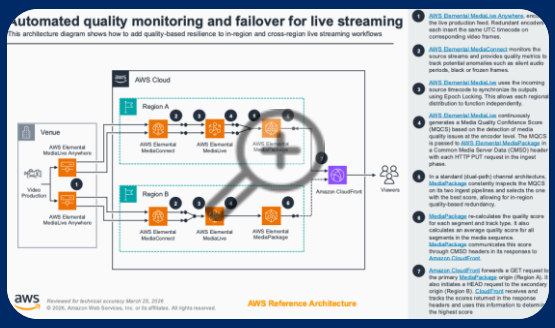

Demo description

Live streaming operations teams struggle to detect and respond to quality issues before viewers are affected. Manual monitoring across multiple availability zones and regions is reactive and resource-intensive, often catching problems too late.

This demonstration shows automated quality monitoring and failover for live streaming pipelines powered by Media Quality-Aware Resiliency (MQAR), an integrated capability between Amazon CloudFront and AWS Media Services that automatically detects, reacts to, and provides visibility into potential issues within live streaming pipelines. Attendees will observe techniques both in-region and cross-AWS Regions, to detect upstream errors and reduce the operational burden of identifying and quickly responding to issues with live streaming workloads running on AWS.

The result: proactive issue detection, automatic recovery without viewer disruption, and reduced operational overhead for live streaming workloads.

Related content

- Amazon CloudFront MQAR Developer Guide
- AWS Elemental MediaLive MQAR User Guide
- AWS Elemental MediaPackage MQAR User Guide
- Blog: Improve your viewers' live streaming experience with Media-Quality Aware Resiliency
- Blog: Build a resilient cross-region live streaming architecture on AWS

Cloud-native video routing at scale

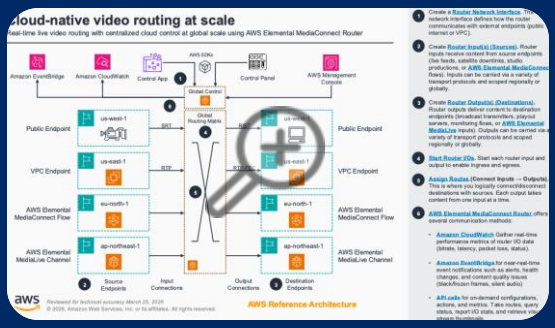

Demo description

Broadcast operations teams managing global live video distribution often struggle with coordination and management in a time sensitive, high-pressure environment. This demo showcases AWS Elemental MediaConnect Router performing real-time live video routing with centralized cloud control at global scale. The demo provides familiar broadcast-oriented operational views from custom React UX frontends controlling MediaConnect Router APIs via the AWS SDK for JavaScript.

Attendees will see cross-region live routing across global infrastructure, familiar management interfaces and broadcast-oriented operational views that provide immediate visibility into stream health and aid troubleshooting. The result: simplified coordination across regions, faster response to routing changes, and scalable operations that reduce complexity in high-pressure broadcast environments.

Related content

- Using the MediaConnect router

Secure streaming delivery using Common Access Tokens and CBOR Web Tokens

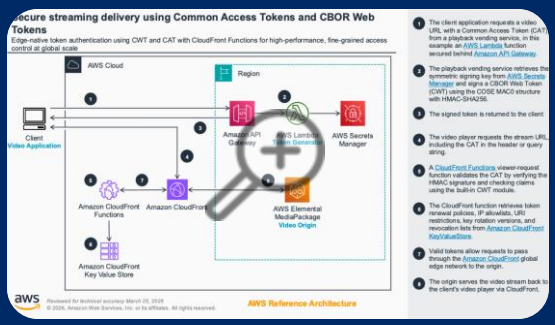

Demo description

Content providers and streaming platforms struggle to secure video delivery while maintaining performance at scale. Traditional authentication methods add latency, create bottlenecks at origin servers, and lack fine-grained access controls for different content types, devices, or geographic regions.

This demonstration shows how CBOR Web Tokens (CWT) and Common Access Tokens (CAT) deliver secure token-based authentication and authorization with CloudFront Functions at CloudFront Edge Locations. CWT provides a compact, binary alternative to JSON Web Tokens (JWT) using Concise Binary Object Representation (CBOR) encoding, while CAT extends CWT with additional fine grained access control including URL patterns, IP restrictions, and HTTP method limitations. Both token types use CBOR Object Signing and Encryption (COSE) for enhanced security and allow developers to implement lightweight, high-performance authentication mechanisms directly at the edge with sub-millisecond execution times

The result: enhanced security without performance degradation, reduced origin server load, and flexible access.

Related content

- CWT support for CloudFront Functions

Create. Connect. Captivate.

Contact us

Learn more