databricks

EXECUTIVE BRIEF

How Media and Entertainment Leaders Are Putting AI to Work

Based on Hundreds of Databricks Deployments

Introduction

Media, entertainment, and advertising companies, from global streamers and broadcasters to media agencies and adtech firms, sit on enormous volumes of operational, customer and behavioral data. This data is normally siloed across teams and various creative assets, locked in different audience platforms, campaign systems, content libraries and ad tech stacks. As a result, decisions that should take minutes take days. Analyst teams are buried in repetitive queries. Campaign and content performance signals that could have been predicted weeks in advance go undetected until they impact engagement, revenue or subscriber churn.

This becomes an even bigger challenge as audiences, the content they consume, and the ways they consume it are growing more fragmented. The media landscape is steadily moving from an attention economy — where success is defined by impressions delivered — to an AI-powered prediction economy, where brands must be able to accurately anticipate audience behavior. The companies pulling ahead aren't simply deploying AI; they're deploying AI on a unified, governed data platform that allows them to move from reacting, to correctly predicting. That's the difference between a pilot that impresses and a system that scales.

Key takeaways

Multi-agent architectures, not single-purpose chatbots, are the production standard for leading companies.

HBO Max powers multiple use cases across content marketing, finance, legal and operations, serving personalized viewing recommendations for 70 million subscribers.

Data governance is the architectural prerequisite for AI at scale.

Unity Catalog's fine-grained access controls, column-level security and full lineage through agent workflows allow media companies to unify subscriber, audience and advertiser data and deploy AI across personalization, content and advertising use cases, without losing control of sensitive data. It also enables privacy-safe collaboration through clean rooms and Delta Sharing while protecting PII, consent and regulatory compliance.

The most effective AI works across every data type — structured and unstructured — in a single coordinated system.

SQL and Genie agents query structured data like subscriber activity, campaign performance and ad inventory stored in Lakebase, while RAG and Vector Search agents analyze unstructured data like content metadata, transcripts, creative assets and campaign briefs. A media planner, content strategist or ad ops lead can ask a question and get an answer that pulls from both real-time operational data and large content or campaign documents, without needing to know where the data lives.

Time to value is measured in weeks, not quarters.

Across Databricks deployments, companies are reaching production faster than expected: Warner Bros. Discovery brings new data into its platform in days, providing a curated selection of programming based on user insights faster than ever before.

The efficiency gains are real and attributable.

Publicis Groupe achieved a 30% productivity improvement across data teams. SEGA accelerated time to insight by 10x with Genie's AI-driven self-service analytics. Comcast Advertising improved data product development speed by 10–30%, enhancing campaign impact. Showtime reduced data pipeline runtimes from over 24 hours to under 4, enabling significantly faster decision-making.

Align your stakeholders before you build.

The most successful deployments involve CDOs, CTOs, media buyers, ad ops and sales leaders, IT security and finance from day one. This cross-functional collaboration is what separates pilots that scale from pilots that stall.

30%	Productivity Increase
$6M	Annual Cost Savings
10x	Faster Time-to-Insight

Use cases

Multi-Agent Supervisor Architectures

Most companies rely on dozens of systems spread across databases and document types. Multi-agent architectures solve this by assigning specialized agents to handle different tasks, all managed by a central supervisor that routes each question to the right place.

HBO Max operates a multi-use-case intelligence platform built on a unified, governed data foundation. By consolidating petabytes of structured and unstructured data into its Databricks lakehouse, the company supports a wide range of critical use cases on a single platform such as personalized content recommendations, real-time customer behavior analytics, content metadata and knowledge graph enrichment, marketing and engagement optimization, subscription and revenue insights, voice/NLP-driven discovery and privacy/legal compliance analysis.

This shared foundation enables cross-functional teams to work from the same source of truth while powering downstream applications simultaneously. With Databricks, HBO Max can deliver real-time analytics, content curation and personalization for nearly 70 million customers, allowing the platform to continuously refine user experiences and drive retention at global scale.

Built on: *Unity Catalog (access controls and data lineage), Genie agents (structured data/SQL), RAG agents (unstructured documents), function-calling agents (vector search), Databricks Model Serving*

Natural Language Data Access and Analytics

What if anyone on your team could get answers from your company's data just by asking a question: no SQL, no analyst queue, no waiting? At SEGA, teams use Genie to query data in plain English and get instant, actionable insights, driving a 10x faster time to insight and true self-serve analytics across the business. Seven West Media uses Genie's conversational interface to get real-time answers on audience, device and content performance in under 25 seconds.

FOX Sports extends natural language data access to fans. The AI-powered semantic search engine built on Databricks lets users ask conversational queries and get fast, context-aware results. Search success rates have more than doubled across hundreds of thousands of requests, delivering more accurate, dynamic content discovery.

Built on: *Databricks Model Serving, open-source LLM support (Llama, Mistral), Genie text-to-SQL, Mosaic AI critique agent framework*

Predictive Audience and Revenue Intelligence

In media, missed engagement and underperforming campaigns are some of the most expensive problems. Predictive models use real-time audience, content and campaign data to forecast behavior before it impacts revenue. By identifying churn risk, predicting content performance and forecasting ad yield and inventory demand, teams can act early. Results include optimizing programming, personalization and monetization strategies before opportunities are lost.

Showtime leverages predictive models to analyze subscriber data, anticipate viewer behavior and optimize content scheduling and programming to boost engagement and reduce churn. Comcast Advertising uses Databricks Apps, packaging predictive models into interactive tools that let business users explore outputs and test different scenarios. This enables teams to adjust ad strategies in real time and make faster, more data-driven decisions across campaigns, sales and customer experience.

Built on: *Databricks streaming pipelines, Delta Lake, real-time ML model serving, Mosaic AI*

Data Governance as Enabler

Governance is what makes enterprise-wide AI deployment possible by giving media companies the confidence to scale across subscriber, audience and advertiser data. Unity Catalog's access controls, data masking and lineage tracking are built into the platform, so AI agents can operate across sensitive datasets (PII, consented first-party data and cross-partner clean room environments), while meeting requirements like GDPR and CCPA to protect consumer privacy and maintain trust. This also helps ensure advertisers can rely on accurate, compliant measurement and attribution.

Condé Nast uses Unity Catalog to securely govern access to its centralized global data, enabling consistent reporting and faster, more reliable analytics. This empowers teams to deliver personalized content experiences at scale across their 37 distinct brands.

Built on: *Unity Catalog (column-level security, dynamic data masking, data lineage tracking through agent workflows, audit logging), fine-grained access controls for multi-agent systems*

Strategic Recommendations

Build Your Foundation

Before your first agent deployment, implement Unity Catalog. Its built-in access controls, data lineage and audit logging are what allow you to expand AI beyond your first team without creating governance debt. Start by auditing high-friction workflows across audience analytics, campaign reporting and content performance; this where teams are still waiting on dashboards, stitching data across platforms or manually analyzing engagement and churn. Flag the highest-volume use cases, like campaign optimization, subscriber insight requests or content performance analysis, as your first automation targets. Then, run a focused two-week pilot on a real problem (e.g., enabling media planners to query campaign performance in natural language with Genie) to prove value quickly.

Scale What Works

Once a workflow is proven in one business unit, treat it as a template for the next. Databricks enables reusable pipelines, shared notebooks and governed data assets that let you take a text-to-SQL workflow or multi-agent architecture and expand it across an organization without rearchitecting. Use MLflow to track model performance and Unity Catalog to manage access as new teams and data domains come online. The goal is to build a library of reusable agent components instead of a collection of isolated deployments.

Skip What Doesn't

Not every investment is worth making. Three to avoid:

- **DIY MLOps infrastructure.** Databricks and major cloud providers have already solved this; building your own means spending engineering time on maintaining infrastructure instead of solving business problems.
- **Single-purpose AI point tools.** Every additional vendor increases the total cost of ownership and adds integration debt.
- **Premature model optimization.** LLMs improve every quarter. Build on an open architecture that lets you swap in better models as they become available without rebuilding your workflows.

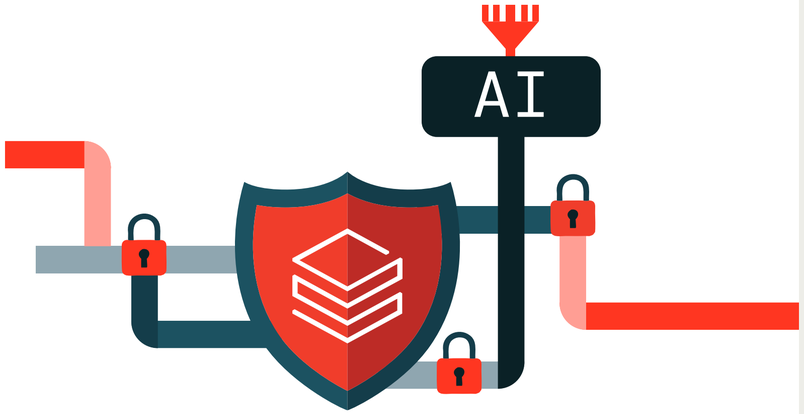

Results from the Field

Three patterns emerge across these deployments:

Efficiency gains are concentrated in data operations: processing, querying and enrichment. This repetitive, high-volume work is exactly what AI handles best.

Speed improvements are transformational, not incremental. With Unity Catalog and generative AI, Epsilon has reduced the time it takes to stand up a data clean room from months to just 48 hours, significantly accelerating client campaign time-to-market.

Scale compounds as each AI workflow absorbs more throughput without adding operational cost. Viacom18 ingests and processes over 45,000 hours of daily content on Voot, generating up to 700GB of data per day, while improving operational efficiency by 26%. With cluster management, auto-scaling and decoupled systems, the platform scales volume without proportional cost increases.

In every case, the underlying enabler is the same: Databricks Platform that connects raw operational data to the AI layer without requiring teams to stitch together point solutions to bridge the gap.

SHOWTIME	6x	Faster Pipeline Processing 24 hours → 4 hours
CONDÉ NAST	$6M	Annual Savings Infrastructure Costs
theTradeDesk	200x	Increase In speed at lower cost
PUBLICIS GROUPE	40–50%	Increase In client campaign revenue

Realistic timeline

When it comes to AI, business decision-makers often wonder: How long before this technology delivers something real?

Based on Databricks deployments across media and entertainment companies:

30 DAYS

Time to Value

Data staff's inboxes at Seven West Media have been reduced by 35% after business users adopted Genie — the chatbot successfully answered 400 unique business questions within the first month

DAYS

Rapid Data Onboarding for Personalization

Warner Bros. Discovery brings new data into its platform in days — not weeks — accelerating content discovery and powering more personalized viewer experiences

2 DAYS

Accelerated Campaign Delivery

Epsilon can set up privacy-safe clean rooms in as little as 48 hours, reducing implementation times from months to days so advertisers deliver more relevant, compliant campaigns faster

SECONDS

Faster Development of Critical Real-Time Insights

SumerSports has cut its time to deliver key post-game AI-powered insights from days to within the minute, allowing football teams to begin strategizing for the next game that much faster.

Real, measurable value can be delivered in weeks, not the multi-year transformation cycles most enterprises are used to. The companies that move fastest have two things in place before they start: a governed data foundation, so agents have clean, accessible data to work with, and a clear business owner who defines what 'working' looks like. Both are achievable before your first sprint begins.

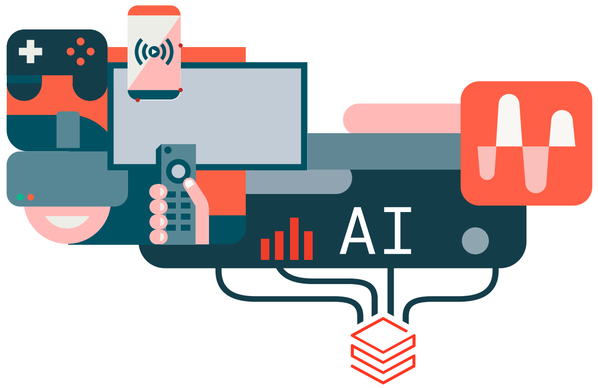

Conclusion

Media and entertainment AI is no longer an experiment; it's an operational capability that leading companies are deploying at scale right now. The companies pulling ahead aren't doing it with a collection of point solutions or a single chatbot. They're doing it with a unified platform that brings data engineering, governed access, multi-agent orchestration and model serving together in one environment.

Databricks is the defensible choice for media AI at scale, and it comes down to three advantages:

Open architecture:
Native support for ChatGPT, Claude, Llama and other open-source models means no model lock-in, and you never have to rebuild when a better model emerges.

Governance at the platform layer:
Unity Catalog makes it possible to deploy AI broadly across an organization without compromising data protection or compliance.

The compounding advantage of a unified platform:
Every new use case built on Databricks shares the same data assets, pipelines and governance framework, so the tenth deployment takes a fraction of the effort the first one did.

The question for media and advertising leaders isn't whether to deploy AI. It's whether to spend the next two years stitching together tools that don't scale, or build on a unified foundation where every deployment makes the next one faster, smarter and more efficient.

databricks