

REACH PLC / MANTIS SOLUTIONS

Video Compliance That Sees What Text-Based Tools Can't

A UK news publisher processes thousands of videos daily. Their brand safety pipeline only understood text. Pegasus on AWS Bedrock closed the gap - cutting compliance review from hours to minutes.

PROCESSING TIME

Hours → Minutes

Per-video compliance review

POC SCOPE

70-80+

Validated with consistent results

ASSET LENGTH

1+ hr Videos

Analyzed as one narrative, not frames

POC → PROD

~6 Months

Q4 2025 → March 2026

The Stakes Just Got Higher



Video is the fastest-growing content format on every news platform and the hardest to govern. Publishers are producing more of it, syndicating more of it, and monetizing more of it through advertising than ever before. That growth isn't slowing down. As AI accelerates content creation across the industry, the volume of video that needs to pass through compliance pipelines will compound every year.

The organizations that win aren't the ones with the most content. They're the ones whose compliance infrastructure can keep pace with it automatically, accurately, and at scale.

Mantis Solutions, Reach PLC's technology arm, saw this problem clearly. **They built the solution before it became a crisis.**

The Problem

Every Publisher Recognizes

Reach PLC is the UK's largest commercial news publisher: the Daily Mirror, the Daily Express, and over 100 regional news brands. Every day, thousands of articles go live across their portfolio. Many carry embedded video. Every piece of that video, if it runs advertising, has to meet brand safety standards.

A video about a natural disaster might be essential journalism. It might also be the wrong context for a luxury travel brand. A segment covering a violent crime is editorially important and commercially sensitive at the same time. The editorial team makes these calls instinctively. The problem is scale: no team is large enough to watch every video before every ad placement.

Before TwelveLabs, Mantis's compliance pipeline was built on text. Articles were analyzed using Natural Language Understanding. Ads were placed or withheld based on what the article said. But video was a blind spot. Manual review couldn't keep pace. Existing tools could label individual frames. They could tell you a frame containing fire, but they couldn't tell you whether the video was covering a wildfire evacuation or showing a building burning down in a crime report.

Mantis didn't need a frame classifier. They needed a model that could watch a video, follow the narrative, and produce a structured brand safety determination with the reasoning to support it.

MEETING THE BRAND SAFETY AND COMPLIANCE BAR

Callum McAdam, Senior Technical Product Manager at Mantis Solutions, led the search for a solution that could meet the compliance bar their advertising clients required at the speed and scale their publishing operation demanded.

"We'd been trying to close the video brand safety and compliance gap for a long time. Frame classifiers and manual review workflows struggled to interpret what was actually happening in the video, whereas Pegasus can actually understand that context. The fact that it explains its reasoning, not just its score, is what makes it workable for our brand safety and compliance team."

— Callum McAdam, Sr. Technical Product Manager, Mantis Solutions

The Automated Brand Safety and Compliance Engine: Pegasus on AWS Bedrock

Mantis deployed TwelveLabs Pegasus on AWS Bedrock to build a fully automated video brand safety and compliance workflow that slots directly into their existing brand safety infrastructure without a new pipeline or a manual handoff.

Pegasus processes the full multimodal signal – visual, audio, and textual – as an integrated whole. It doesn't sample frames and guess. It tracks events, context, and intent across the entire duration of a video.

A compliance analyst doesn't get a list of flagged timestamps to scrub through. They get a structured assessment with a pass/fail determination, a safety score, and the reasoning that produced it, plus thumbnail previews at flagged moments. No one has to re-watch an entire video to understand why it was flagged.

THREE CAPABILITY PILLARS

01

Video-to-article generation

Pegasus extracts structured metadata from video content, enabling Mantis to treat video as a first-class content object in their pipeline.

02

Brand safety and contextual analysis

Pegasus evaluates video against Mantis's compliance taxonomy and produces structured safety assessments that feed directly into their core brand safety infrastructure.

03

Search and recommendations

Compliance outputs combined with video-level understanding power content discovery and contextual ad targeting, turning compliance infrastructure into a revenue-generating product.

Specificity and Structured Reasoning

Mantis didn't run a demo. They ran Pegasus against their actual content: the edge cases that consistently broke other tools — long-form news segments, emotionally ambiguous footage, videos where the visuals and the narration tell different stories.

The results delivered a level of specificity that prior approaches couldn't reach. Moving beyond simple object detection, the model provided structured reasoning, detailing exactly what was happening in the video, why it met or failed compliance thresholds, and the confidence score behind every determination.

That accuracy and structured output format were the reasons the engagement moved to production.

The engagement began in Q4 2025. Mantis validated Pegasus against 70–80+ videos during the POC with consistently positive results on brand safety accuracy. The deal closed and moved to production in March 2026 — approximately six months from first engagement to live deployment.

PROMPT ENGINEERING FOR COMPLIANCE

Mantis' data science team worked with TwelveLabs field engineering to optimize Pegasus prompts for their specific compliance taxonomy. For longer content, the team developed a multi-level prompting strategy: breaking segments down, analyzing each, then synthesizing a unified compliance determination across the full asset. Narrative context preserved. No frame-level information loss.

AWS PARTNERSHIP

Pegasus was deployed through AWS Bedrock, giving Mantis enterprise-grade infrastructure with familiar security controls, billing, and governance without the operational overhead of self-hosting a video foundation model. The deal was executed through the AWS Marketplace.

Mantis is now onboarding additional content volumes and expanding into contextual ad targeting powered by video-level understanding. They have also secured Pegasus 1.5 beta access and are sharing labeled production data with TwelveLabs to support compliance-specific improvements that benefit the broader ecosystem.

Delivering Speed and Business Value

Mantis's operations team no longer watches footage to make brand safety and compliance calls. Pegasus watches it for them and explains its reasoning.

- 01

Brand safety and compliance processing: Hours to minutes. Manual review workflows have been replaced by automated analysis that runs at pipeline speed.
- 02

Full-length video analysis. Assets exceeding one hour are analyzed as complete units, preserving the narrative context that frame-by-frame approaches lose entirely.
- 03

Structured outputs, no manual handoff. Automated pass/fail determinations feed directly into the existing brand safety pipeline, delivering both the final assessment and the supporting reasoning.
- 04

Three-tier product architecture. Mantis can now package video intelligence for advertising clients, transforming compliance infrastructure into a revenue-generating asset.
- 05

Pegasus 1.5 beta access secured. Capabilities continue to expand in production as the engagement evolves.

A Repeatable Pattern for the Industry

The engagement with Mantis Solutions represents a repeatable pattern for the industry. Any organization monetizing video through advertising faces the same critical gap: text-based brand safety tools cannot see the content they govern. Pegasus on Amazon Bedrock closes this gap with a production-ready solution that integrates seamlessly into existing enterprise infrastructure.

By sharing labeled production data to support the Pegasus 1.5 beta, Mantis is driving brand safety and compliance-specific improvements that benefit the entire ecosystem. Real data, real requirements, and real accountability are how these models evolve to meet the next generation of publishing challenges.

That shift freed the team to focus on the edge cases that actually require human judgment, rather than spending hours on content that clearly passes or clearly fails. As Mantis' video volume grows, the system scales with it. The organizations building this compliance infrastructure now won't just be more efficient; they'll be operating in a fundamentally different category from publishers still relying on manual review and frame-level classifiers as video volume compounds around them.

TECHNICAL ARCHITECTURE

TwelveLabs model	 Pegasus (video language model)
Deployment	AWS Bedrock
Procurement	AWS Marketplace
Architecture	3-tier (compliance, metadata, and recommendations)
Content scope	Long-form news video, 1+ hour assets
Prompt strategy	Multi-level prompting for long-form content
Output format	Structured (safety score, pass/fail, reasoning, and timestamped thumbnails)

About the Companies



TwelveLabs builds video foundation models that transform how machines understand video. Its multimodal AI platform enables enterprises to search, analyze, and generate insights from video at scale. Based in San Francisco and Seoul.



Reach PLC is the UK's largest commercial news publisher, home to the Daily Mirror, Daily Express, and 100+ regional brands. Mantis Solutions is Reach's technology division providing brand safety, contextual targeting, and data infrastructure for digital advertising.

mantissolutions.com

Get in Touch

- twelvelabs.io
- sales@twelvelabs.io
- docs.twelvelabs.io

