ViewSignals: Inspection on AWS

aws partner network

Executive Summary

ViewSignals was created to solve one of underwriting's biggest timing problems. Traditional field inspections often arrive weeks after a policy decision has already been made. Working with Hex Networks, ViewSpection built an AI-powered inspection intelligence layer that analyzes ViewSpection self-inspection photographs, prior inspections, and property intelligence within minutes—bringing actionable risk signals to underwriters during submission and pre-bind workflows. The resulting product, ViewSignals, enhances ViewSpection's self-inspection platform by using AI to surface structured risk signals from self-inspection imagery, historical inspections, and property intelligence so underwriters can review submissions faster without removing human judgment.

Business Challenge

In most underwriting workflows, the property inspection arrives last - often weeks after a policy is bound. Underwriters select and price risk with limited visibility into the property, then absorb whatever the inspection later reveals. Reviewing that evidence by hand is slow and inconsistent from one reviewer to the next.

Viewspection needed more than a simple image classifier. They required a system that could reason about a property the way an experienced inspector does, quantify how confident it was in each finding, automatically handle the straightforward cases, and reliably escalate the ambiguous or high-risk ones to a person. It also had to protect privacy, control cost, and scale elastically with demand.

Solution Provided by Hex Networks

Hex Networks partnered with Viewspection to design and deliver a fully serverless property inspection analysis platform on AWS. The solution orchestrates a suite of AWS Lambda functions with AWS Step Functions, using Amazon Rekognition for image quality and object detection and Amazon Bedrock for deep property analysis with per-finding confidence scoring.

Rather than a single model call, the platform is a coordinated pipeline: images are validated and quality-filtered in parallel, analyzed for damage and risk, scored, and passed through a multi-check decision gate that determines whether an inspection can be auto-approved or must go to a human reviewer. Prior reports and image findings are indexed together in a property-scoped vector store, giving every analysis the full historical context of the property. The entire environment is defined as code in Terraform and deploys repeatably across accounts and regions.

About Viewspection

ViewSpection

Viewspection is a property inspection technology company serving the insurance industry. Its platform, ViewSignals, brings inspection intelligence forward in the underwriting workflow - surfacing risk signals from imagery, prior inspections, and property data at submission and pre-bind. Precision, trust, and human judgment are central to everything Viewspection delivers.

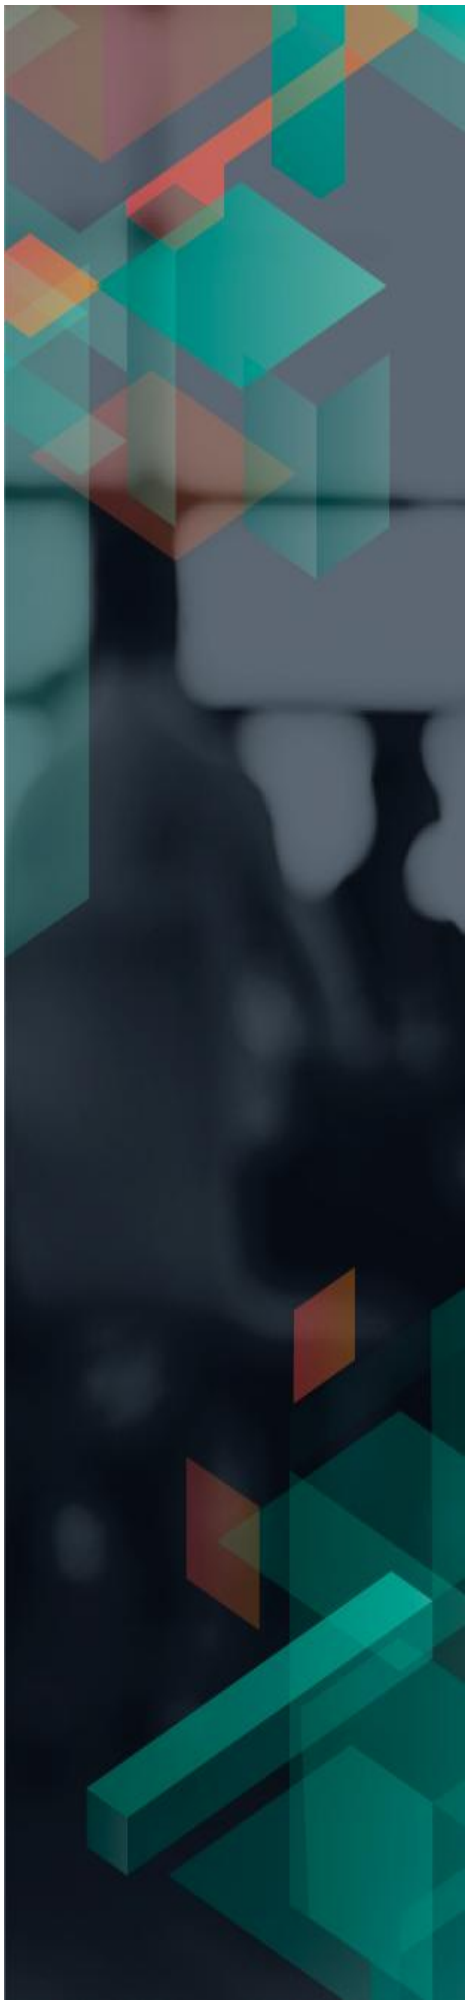

Intelligent Image Analysis Pipeline

Each uploaded photograph flows through a parallel quality-and-safety stage before analysis. Amazon Rekognition scores every image for sharpness, brightness, relevance, and resolution, while dedicated detection steps flag people or personally identifiable content and surface obvious risk indicators. Low-quality or irrelevant images are caught early, so the AI spends its effort only on usable evidence.

Qualifying images are then analyzed by Amazon Bedrock, which produces a detailed property assessment - condition, identified issues, risk factors, and underwriting considerations - and assigns each individual finding a confidence score from 0 to 100. Findings that fall below a configurable threshold are suppressed, so noisy “maybes” never reach the final report.

The Decision Gate: Confidence-Based Automation

At the heart of the platform is a decision gate that evaluates every inspection against six checks: confidence, image quality, internal consistency, risk severity, privacy, and insurability. Per-finding scores are rolled up into a single inspection-level confidence number and a HIGH / MEDIUM / LOW certainty band, giving reviewers and downstream systems both granular and at-a-glance signals.

When all six checks pass, the inspection is approved automatically and the carrier is notified through Amazon SNS. If any check fails, the inspection is routed to a QA queue on Amazon SQS and tagged HIGH, MEDIUM, or LOW priority based on which checks failed - so reviewers naturally pick up the most urgent cases first. A configurable mode lets Viewspection run fully automated or send every inspection to a human, with no code change.

Automated Report Processing

Prior inspection reports rarely arrive as clean data - they are PDFs. The platform ingests these automatically, extracts their text and tables with Amazon Textract, and breaks the content into semantically meaningful chunks.

Each chunk is embedded with Amazon Titan Embeddings and indexed for retrieval, turning static documents into a searchable knowledge base for the property. Historical findings become instantly available to inform new inspections and to detect discrepancies between what a report claimed and what the images now show.

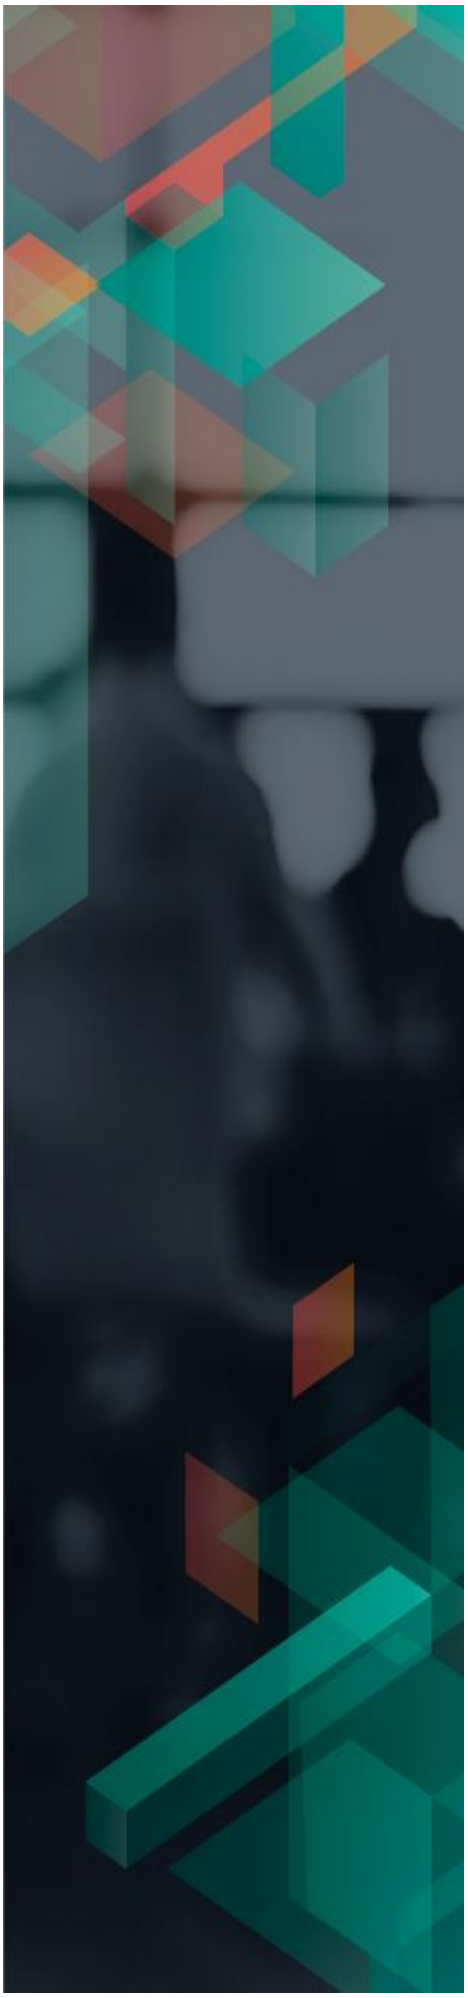

Unified RAG: Retrieval Across Reports and Images

A defining feature of ViewSignals is that it treats images and documents as one body of knowledge. The AI analysis of every image is converted into text, embedded with Amazon Titan Embeddings, and indexed into the same Amazon OpenSearch Serverless vector store that holds the report chunks.

Because both sources share a single property-scoped index, any retrieval for a property surfaces both image-derived and document-derived context. Every subsequent inspection of that property automatically benefits from the accumulated history of all previously processed images and reports - the system grows more informed about each property over time.

A Serverless Architecture Built for Scale and Cost

The entire platform is serverless. AWS Lambda and AWS Step Functions run the processing pipeline, Amazon S3 stores media and outputs, Amazon DynamoDB holds inspection state, and Amazon API Gateway exposes the submission and review APIs. There are no servers to patch or scale; the system scales elastically with inspection volume and costs nothing when idle.

Cost is treated as a first-class concern. Every Bedrock and Titan call records its token usage, which is priced and written back onto the inspection record as a stage-by-stage cost breakdown. Viewspection gets a precise per-inspection dollar figure, making it straightforward to price the service to carriers and to confirm that model or prompt changes actually reduce cost.

Security, Privacy, and Human Oversight

Privacy is enforced in the pipeline itself. Dedicated detection flags faces and personally identifiable information, and any inspection touching privacy-sensitive content is automatically routed for human review rather than auto-approved. Access to AWS resources is governed by least-privilege IAM, and all processing runs inside a controlled AWS account.

For the cases that need judgment, a review dashboard gives experts a clear queue ordered by priority. When a reviewer corrects or confirms an AI result, both the original output and the human decision are captured - building an ever-growing, labelled dataset of ground truth that feeds future evaluation and model improvement.

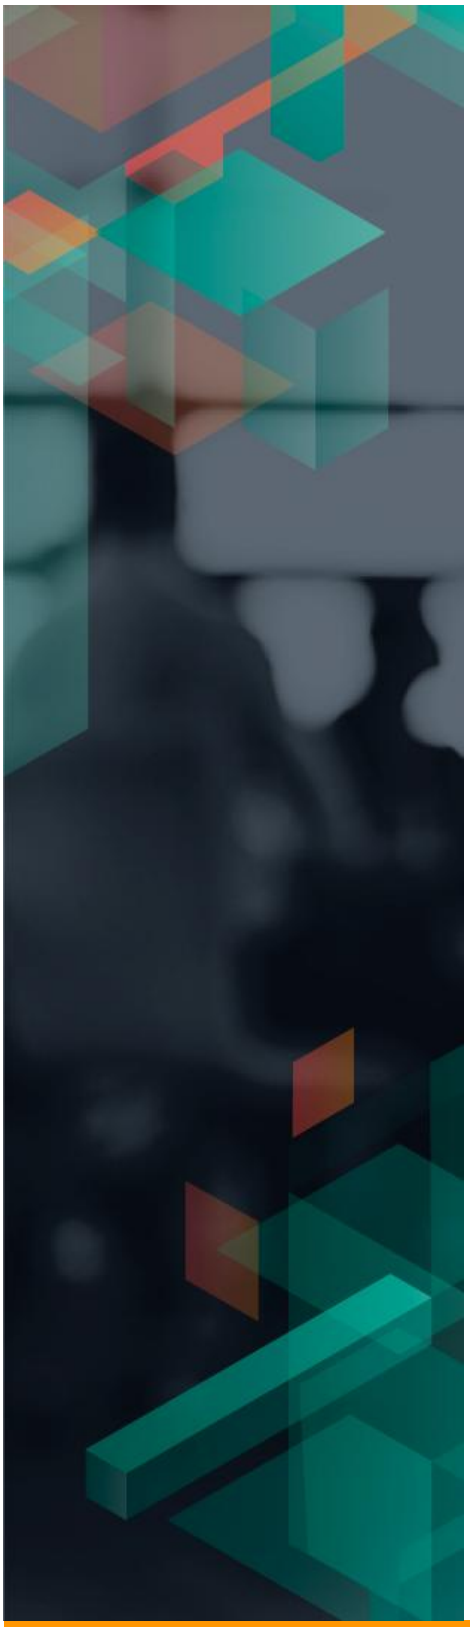

Results and Impact

ViewSignals transforms inspection review from a manual bottleneck into an automated, scalable service. Straightforward inspections are analyzed and approved in minutes without human effort, while ambiguous or high-risk cases are reliably escalated with clear priority - letting expert reviewers focus where they add the most value. The goal is not to replace underwriting judgment - it is to bring meaningful property intelligence forward so better underwriting decisions can be made earlier in the workflow.

Potential benefits include minutes instead of days, confidence scoring, structured outputs, consistent review, prioritized human review, scalable processing, measurable AI cost, and reusable inspection knowledge.

Why AWS

Building on AWS let Hex Networks assemble advanced capabilities - generative AI with Amazon Bedrock, managed vector search with Amazon OpenSearch Serverless, document intelligence with Amazon Textract, and image analysis with Amazon Rekognition - without operating any of the underlying infrastructure. Orchestration with AWS Step Functions and AWS Lambda tied these services into a single, resilient workflow.

The result is a production-ready platform that is secure, elastic, and cost-transparent by design. As Viewspection grows, the same architecture extends naturally to new carriers, higher volumes, and richer models - a durable foundation for the future of AI-assisted property inspection. This collaboration highlights Hex Networks' expertise in delivering not just technical solutions, but production systems that clients can trust and scale.

About Hex Networks

Hex Networks is focused on creating reliable and secure environments to solve issues quickly and efficiently. We identify areas of opportunity in IT systems to manage overall efficiency while reducing costs by implementing tested and reliable solutions. We cater to our customers by building personal relationships for effective problem resolution. Unlike other