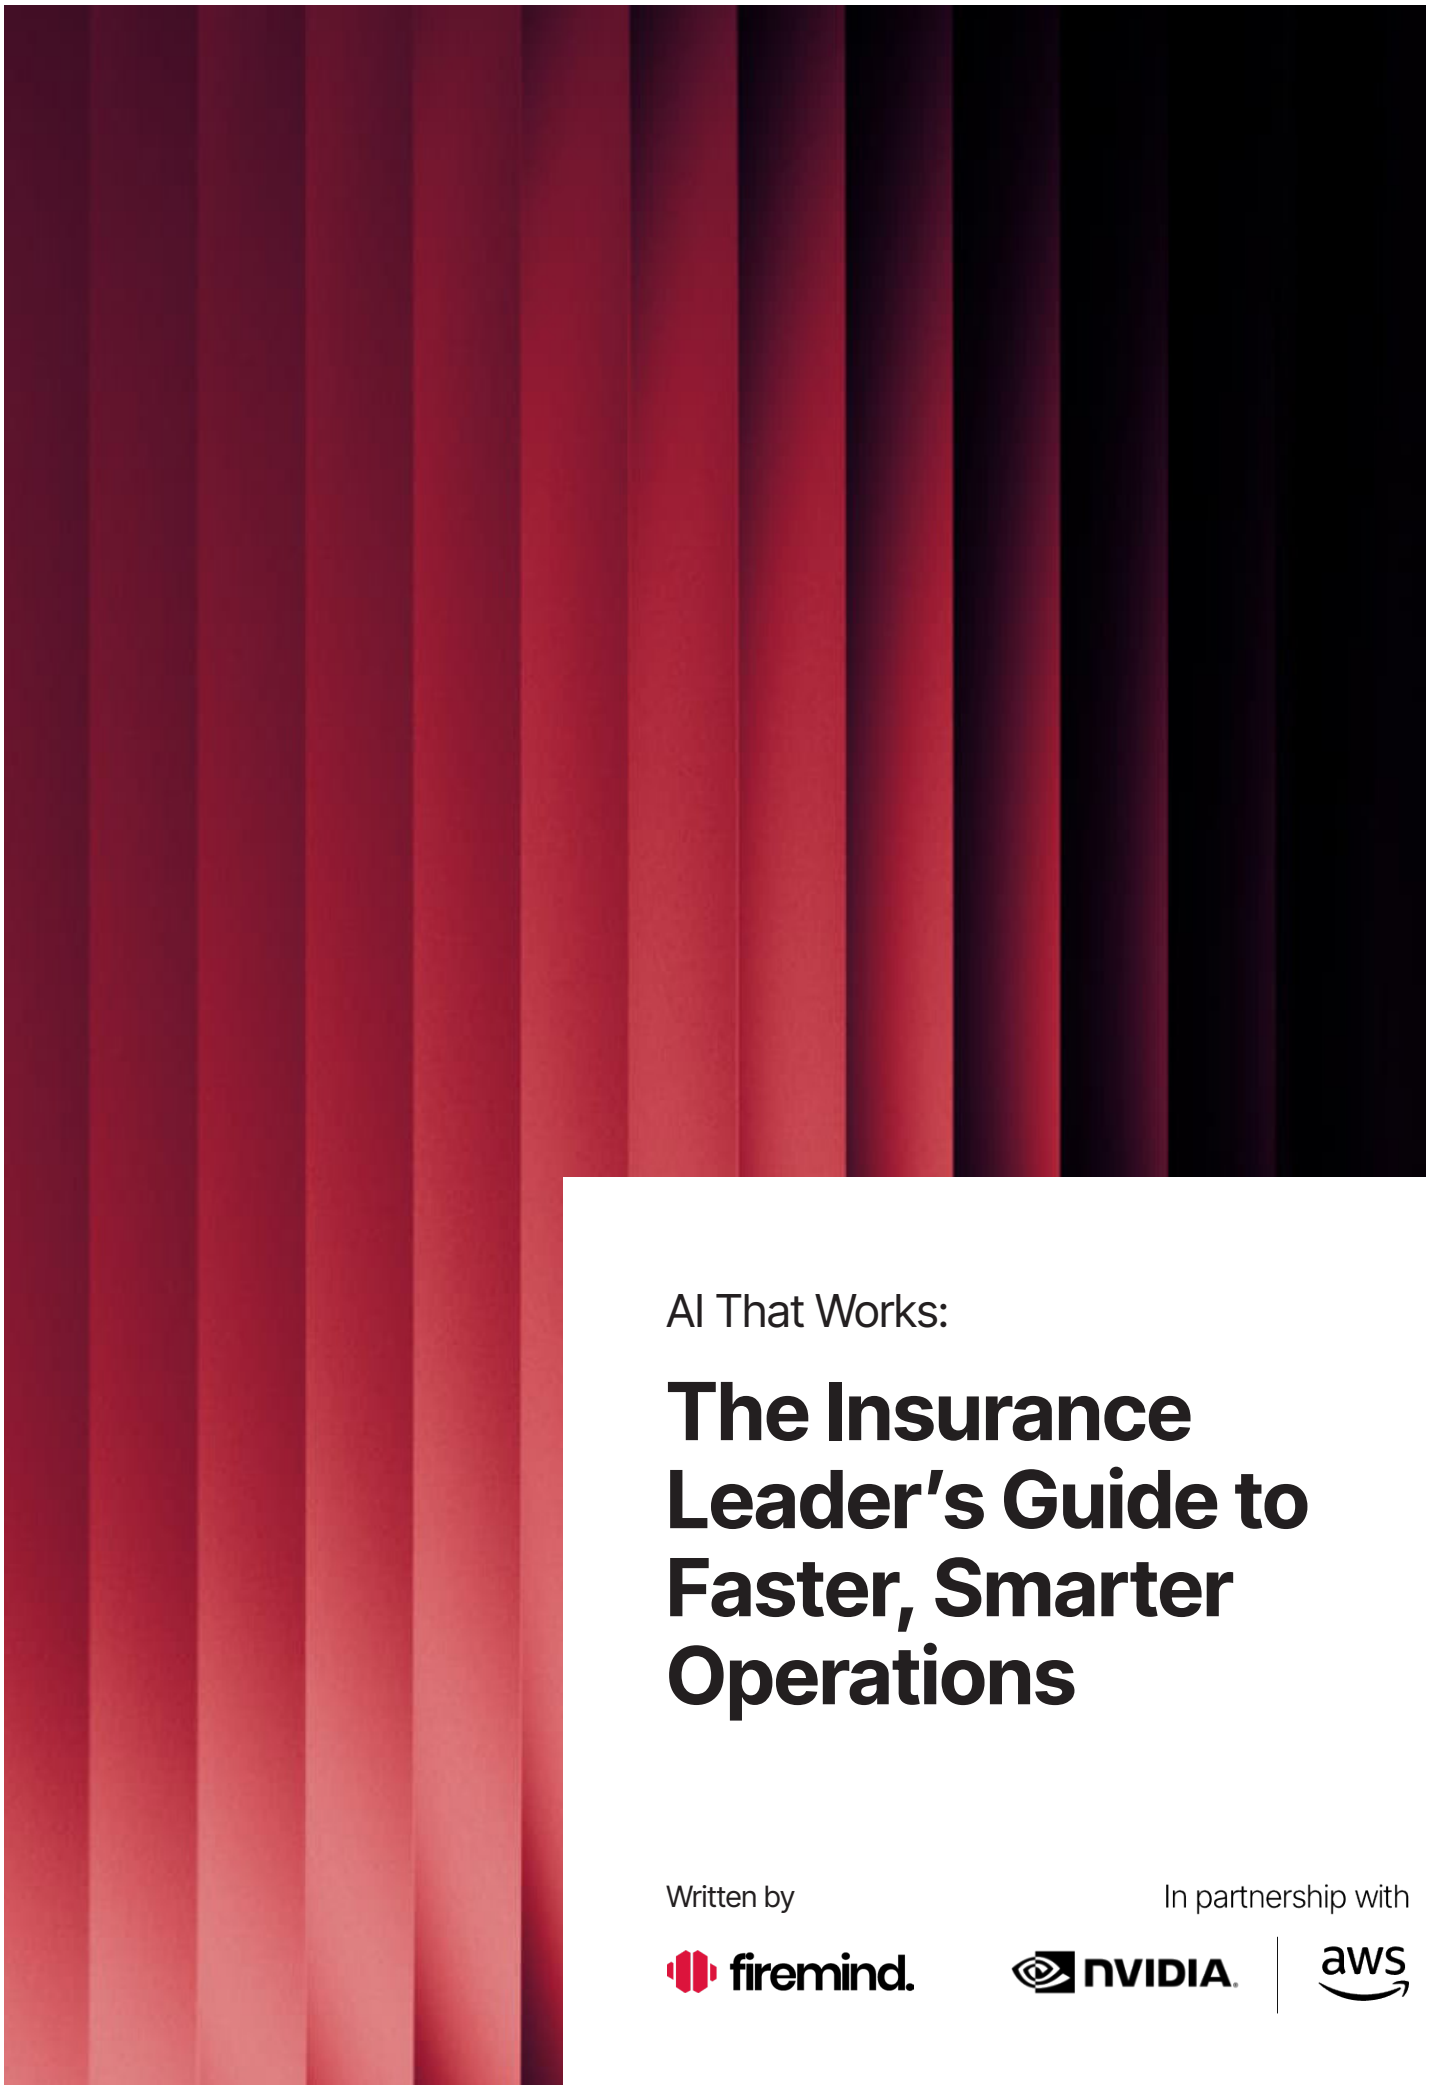

AI That Works:

The Insurance Leader's Guide to Faster, Smarter Operations

Written by

firemind.

In partnership with

NVIDIA.

aws

Executive summary

Insurance operations face rising pressure from all sides: increasing submission volumes, more complex claims, higher regulatory expectations, and the need to control expenses. Leaders across the industry are asking the same question: How can we deliver faster, better outcomes without increasing headcount or rebuilding core systems?

AI now offers a practical answer. Modern AI reads documents in minutes, summarises submissions, structures evidence, and highlights key insights. It expands capacity without expanding teams. It improves consistency and makes decisions easier to audit.

AI has moved beyond experimentation. It is now a foundation of high-performance insurance operations.

The organisations moving fastest are not pursuing large, multi-year transformations. They start with one workflow, deliver value within weeks, and scale what works. This guide shows how any insurer can follow that path.

CHAPTER 1

Why AI Has Become an Operational Priority

AI's value became clear as soon as it touched real insurance workflows. Underwriting teams, claims handlers, and assistance specialists realised the same thing: the work is not slowed by judgment; it is slowed by the time required to gather and interpret information.

Insurance is built on documents, context, and expert evaluation. Before any decision is made, teams must read, understand, and organise large amounts of unstructured information. AI accelerates this preparatory work.

AI enables:

- Faster access to relevant information
- More consistent outcomes across teams and regions
- Clearer audit trails
- Improved governance
- Reduced cognitive load

AI does not replace expertise. It increases the amount of time experts can spend applying it.

For the first time, insurers can meaningfully increase operational velocity without replacing legacy systems. This makes AI central to the modern operating strategy.

CHAPTER 2

The Market Forces Reshaping Insurance Operations

Insurance leaders are navigating forces that traditional processes cannot absorb.

- Customers expect rapid responses.
- Brokers want faster underwriting clarity.
- Regulators require transparency and fairness.
- Claims are becoming more intricate and data heavy.
- Underwriting submissions continue to increase.
- Legacy systems remain essential but limiting.

These pressures make operational performance a competitive advantage. AI addresses the gap by giving teams structured insight at the speed of business.

CHAPTER 3

The Operational Friction Slowing Down Performance

Insurance organisations face a consistent set of friction points, all rooted in manual, document-driven workflows.

Typical sources of slowdown include:

- Manual document processing that absorbs expert time
- Slow triage making it hard to prioritise high-value work
- Inconsistent decisions due to subjective interpretation
- Limited capacity with teams overloaded
- Data silos making information hard to compile
- Governance concerns slowing adoption

As volume grows, these pressures compound. Backlogs form, accuracy varies, and costs rise. AI directly addresses these challenges by handling the preparation work so experts can focus on judgment.

CHAPTER 4

The AI Capabilities Delivering Real Impact Today

AI is now producing meaningful results in everyday workflows, but the biggest gains come from using AI in places where unstructured information meets operational pressure.

Underwriting: AI as a Preparation Engine

AI quickly digests submissions, extracts attributes, identifies unusual elements, and produces structured summaries. This gives underwriters immediate clarity about appetite, exposure, and next steps.

More importantly, AI reduces noise. Underwriters no longer sift through dozens of documents or inconsistent formats. They begin each case with a clean, consistent view of the risk.

The insight upgrade: Underwriting moves from “reactive document review” to “quick intelligence gathering.” This is the real step-change.

Claims: AI as a Consistency Layer

AI organises evidence, highlights key case details, identifies missing information, and surfaces similar historic claims. This reduces rework and supports fair, consistent outcomes.

The deeper insight is that AI standardises early-stage claims handling. By front-loading structure, it becomes easier to spot outliers, confirm coverage, or move quickly on simple cases.

Assistance: AI as a Clarity Tool in Complex Moments

Assistance teams often operate under high stress and variable formats. AI summarises medical reports, interprets case notes, and provides a consistent baseline for decisions such as treatment options or repatriation.

The real transformation is not speed. It is confidence. Teams can rely on AI to convert dense medical language into actionable context.

Cross-Workflow Impact

Across all three areas, AI acts as the same thing: a clarity engine that removes noise so people can act quickly and confidently.

CHAPTER 5

What Real Results Look Like

AI-enabled workflows drive improvements that go beyond efficiency. They reshape the operating rhythm of entire teams.

1. Speed with stability

The biggest win is not just faster processing. It is the stability that comes with it. Teams stop firefighting and start operating proactively. Underwriters respond faster. Claims move forward cleanly. Assistance cases progress with more certainty.

2. Higher-quality decisions

Because AI provides consistent preparatory context, decision quality improves. Experts are no longer limited by how much documentation they can manually process in a day.

3. Reduced operational drag

Meeting prep, reading packs, and first-pass assessments shrink dramatically. The hidden cost of “context-building” goes down. Teams feel lighter, not just faster.

4. Stronger governance

Auditability improves because decisions start from standardised summaries and consistent logic. This is becoming essential in markets facing increased regulatory attention.

5. Measurable value

Examples from insurers deploying AI workflows include:

- Document review time reduced by 75%
- Underwriting cycles accelerated 40–60%
- Team capacity increased by 2–3×
- Operational cost per workflow reduced 20–30%

The deeper insight: AI doesn’t just speed up the work. It changes how the work is done. It enables a more modern, predictable, and resilient operating model.

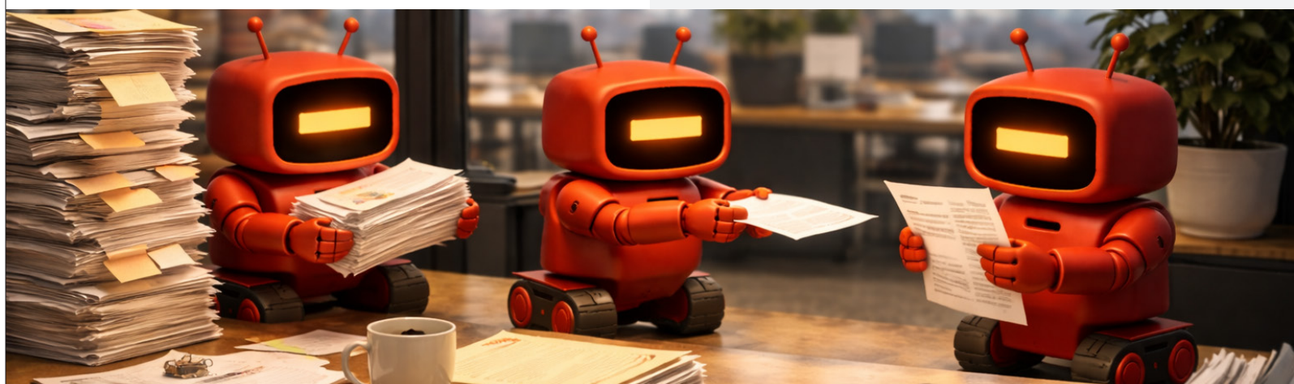

CHAPTER 6

Real Insurance Case Studies

Case Study 1 Carrier A: Accelerated Underwriting

A specialty carrier struggled with rising submission volume. Underwriters were losing time to document-heavy preparation work, making it harder to respond quickly to brokers and maintain appetite discipline.

What AI enabled

- Automated summarisation of submission packs
- Extraction of key risk attributes and anomalies
- Prioritisation of in-appetite risks
- Consistent first-pass briefing packs for underwriters

Impact (anonymised and representative)

- Underwriting preparation time reduced by 50–60%
- Meaningfully more submissions reviewed each day
- Faster broker responses and fewer missed opportunities
- Clearer documentation supporting governance and audit

Case Study 2 Broker B: Intelligent Document Processing at Scale

A large commercial broker processed thousands of documents each week across placement, renewals, and client onboarding. Manual extraction slowed placements and introduced inconsistencies.

What AI enabled

- Automated extraction across all inbound formats
- Normalisation of inconsistent documents
- Entity matching and validation
- Clean, structured outputs ready for immediate placement use

Impact (anonymised and representative)

- Document preparation time reduced by over 90%
- Significant reduction in rework due to data inconsistencies
- Faster placement cycles and improved internal throughput
- Teams able to focus on client strategy, not admin

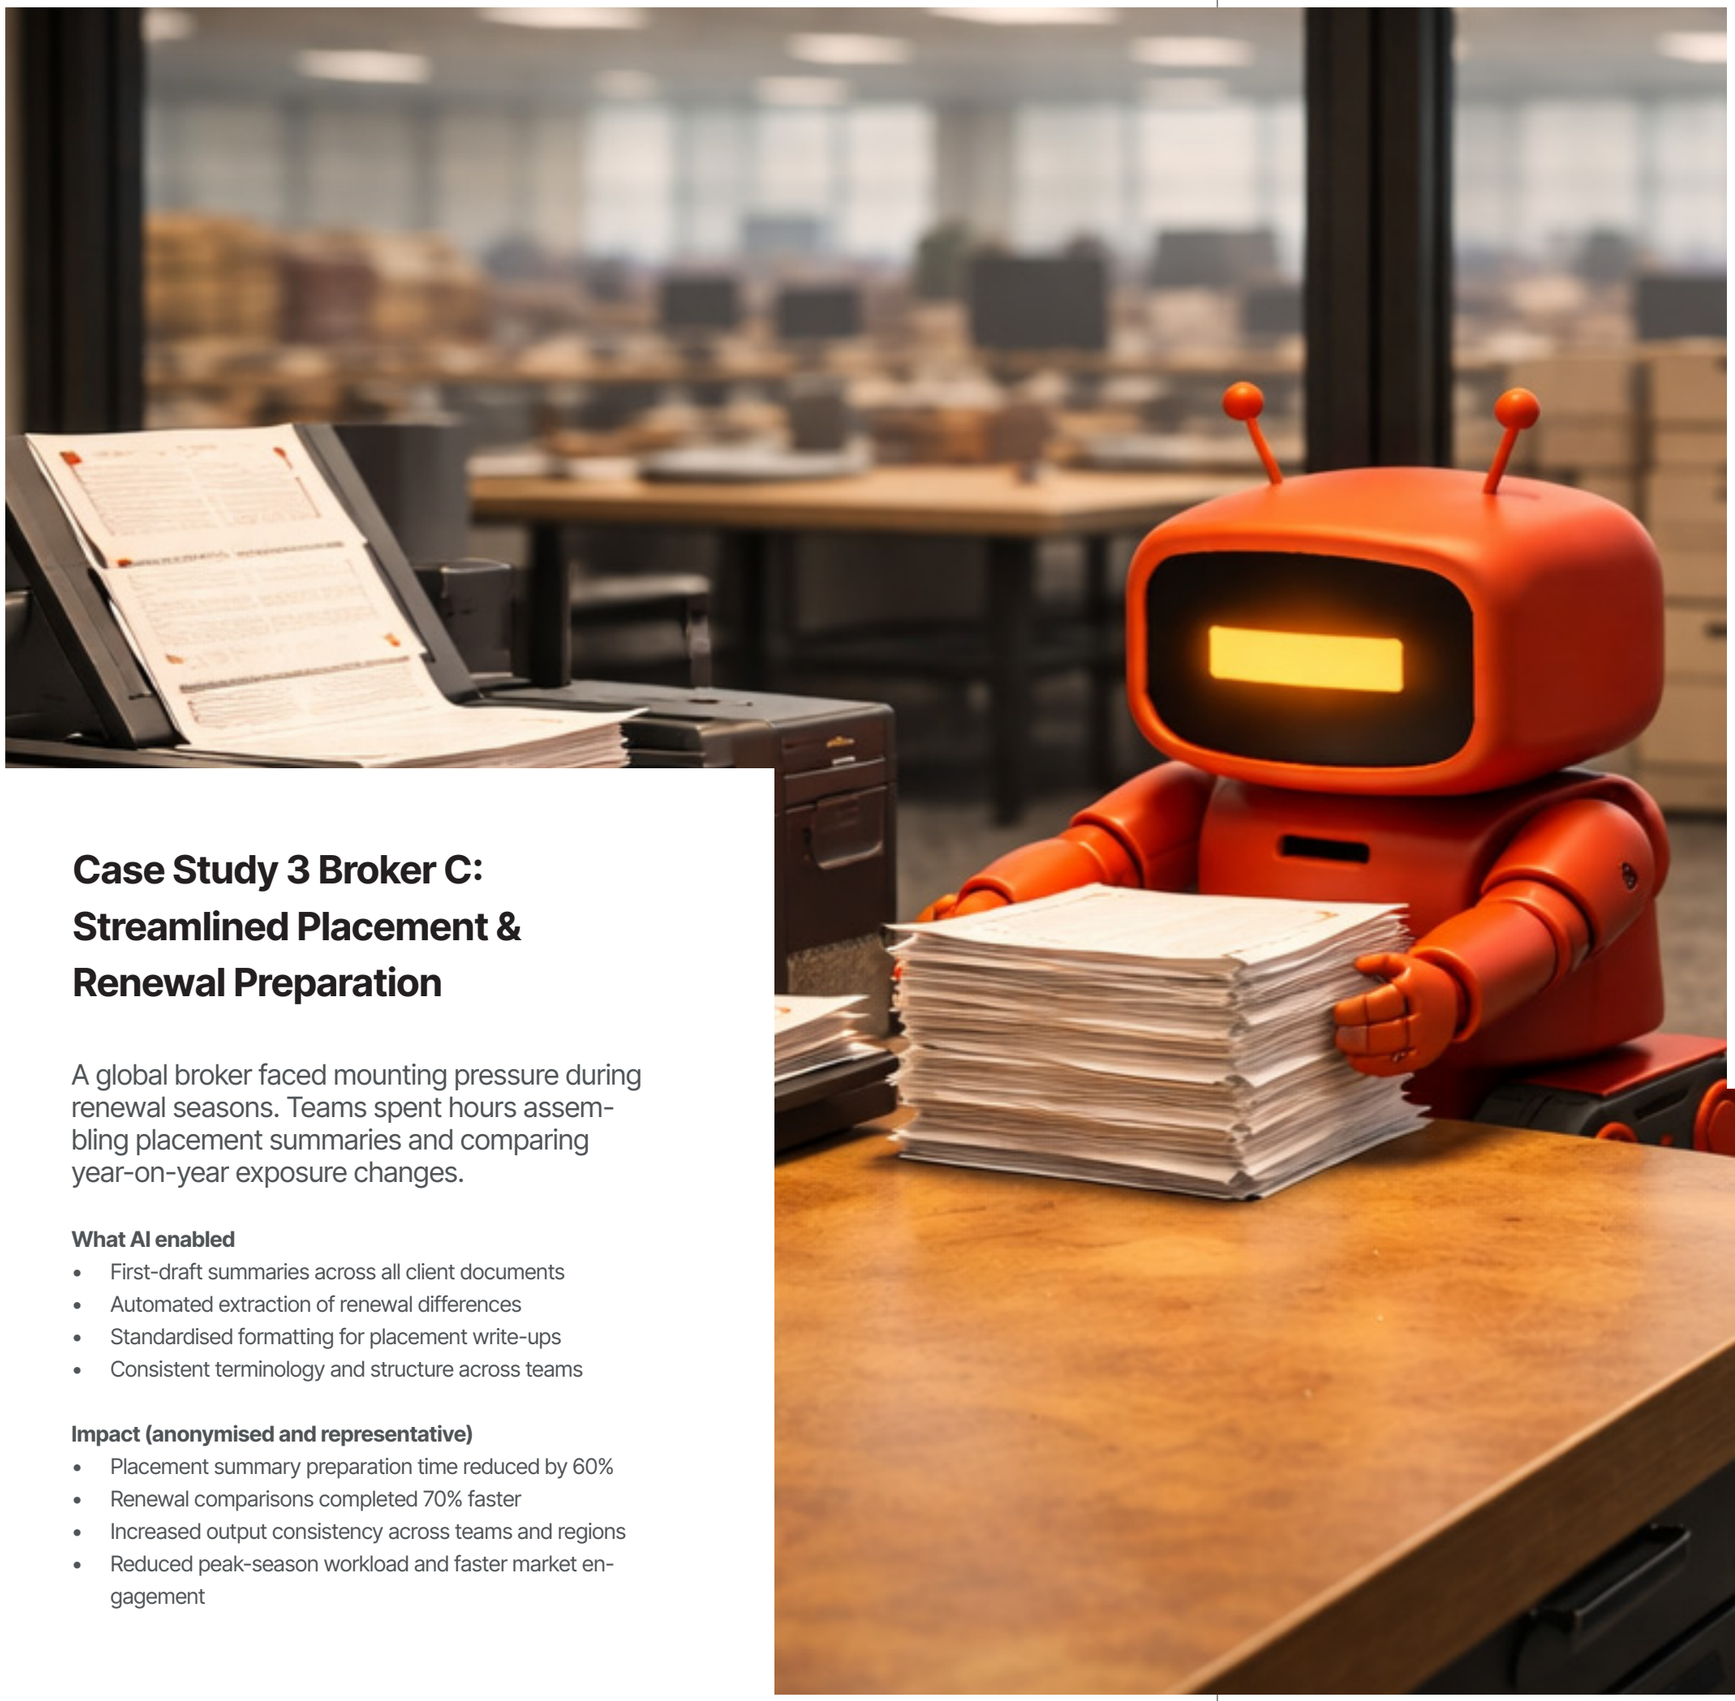

Case Study 3 Broker C: Streamlined Placement & Renewal Preparation

A global broker faced mounting pressure during renewal seasons. Teams spent hours assembling placement summaries and comparing year-on-year exposure changes.

What AI enabled

- First-draft summaries across all client documents
- Automated extraction of renewal differences
- Standardised formatting for placement write-ups
- Consistent terminology and structure across teams

Impact (anonymised and representative)

- Placement summary preparation time reduced by 60%
- Renewal comparisons completed 70% faster
- Increased output consistency across teams and regions
- Reduced peak-season workload and faster market engagement

Case Study 4 Aggregator D: Improving Customer Conversion with AI Agents

A major price comparison platform needed to help customers understand why insurance prices vary. Lack of transparency caused drop-off and reduced trust.

What AI enabled

- A multi-agent orchestration framework to analyse quotes
- Retrieval-augmented generation for safe internal data access
- Real-time integration with MOT records, profiles, and market benchmarks
- Personalised insurance explanations delivered in natural language
- Sub-30-second end-to-end response times

Impact (anonymised and representative)

- 90% accuracy in multi-source data analysis
- Up to 30% faster insight generation
- 20–30 second latency for full analysis and LLM reasoning
- More trusted, transparent price explanations
- Higher conversion through clear, contextual recommendations

CHAPTER 7

What Becomes Standard Next (with AWS, NVIDIA, Firemind)

Insurance operations are entering a new phase. The industry is moving toward AI-enabled workflows that handle high-volume information and support decision-making with clarity and speed. Over the next few years, several developments will define the new standard for how underwriting, claims, and assistance teams operate.

AI-First Workflows Become Normal
Document intelligence, email triage, summarisation, and automated preparation will become the default way insurance teams handle information. Teams will begin their work with structured insight instead of raw documents. The pace of operations will increase as background tasks are handled automatically.

The AWS Underwriting Workbench Becomes a Core Platform
The AWS Underwriting Workbench reflects a shift toward unified workspaces designed around underwriting workflows. It brings together documents, enrichment data, actuarial insights, workflow status, and integrated collaboration.

Underwriters work in a single environment where AI provides instant summaries and context. The workbench becomes the command centre for underwriting decisions and replaces the patchwork of shared drives, inboxes, and spreadsheets that slow teams down today.

AI capabilities integrate directly into the workbench, giving underwriters a modern platform designed for speed, governance, and consistency.

NVIDIA Expands the Technical Possibilities for Insurance AI
NVIDIA's technology ecosystem is becoming increasingly relevant to insurers who want more control over how AI works inside their organisation. NVIDIA brings:

- Advanced model optimisation for insurers with specific accuracy or performance needs
- Hardware acceleration that supports large-scale document processing
- Tooling that allows organisations to customise models for niche or high-stakes lines
- Secure, enterprise-focused AI infrastructure

For insurers exploring deeper AI adoption, including retrieval-augmented generation, custom model training, or low-latency processing, NVIDIA's technology stack opens new possibilities that complement cloud-native services. One example of how organisations are addressing these challenges in practice is the Firemind Blueprint for NVIDIA AI on AWS.

Firemind Blueprint for NVIDIA AI on AWS

Overview
The Firemind Blueprint for NVIDIA AI on AWS enables enterprises to build and scale advanced AI workloads while retaining full control over their data, models, and intellectual property.

Customer Challenge
Organisations want to leverage NVIDIA-accelerated AI on AWS, but face challenges around operational complexity, governance, and long-term ownership when using fully managed or SaaS-based solutions.

Solution
Firemind delivers a production-ready AI platform that combines AWS-native services, NVIDIA-optimised infrastructure, and the NVIDIA AI Enterprise software stack, with a fully managed operating model. This approach provides the flexibility and performance of a self-managed stack, without the burden of day-to-day operations.

At the core is a ready-to-use AI Sandbox, giving data scientists and AI teams a governed environment to experiment safely with NVIDIA-accelerated AI on AWS. For a fixed entry investment, customers receive a secure, compliant sandbox with GPU-backed compute, NVIDIA AI Enterprise tooling, and preconfigured governance enabling teams to move from experimentation to production-ready architectures without compromising control, security, or sovereignty.

Business Outcomes

- Faster time to value for AI initiatives
- Predictable performance and scalability at enterprise level
- Full ownership and control of data, models, and IP

CHAPTER 8

Firemind Becomes the Operational Glue That Turns Technology into Outcomes

Technology alone does not deliver value.
Insurance leaders succeed when AI becomes operationalised, when it fits into real workflows, strengthens governance, and improves decision-making.

This is where partners like Firemind play a critical role.

Firemind specialises in industries that depend heavily on documents, compliance, and judgement. The team focuses on turning AI into daily operational tools that serve underwriting, claims, assistance, and shared services.

Firemind brings:

- Deep insurance domain knowledge
- Reusable workflow patterns proven across carriers and brokers
- Governance frameworks aligned to regulation
- Experience integrating AI into legacy systems
- Alignment with AWS and NVIDIA's technology ecosystems
- The ability to teach internal teams how to use, supervise, and trust AI

Firemind's role is simple to explain but difficult to replicate: turn AI into something that actually works in operations.

Insurers avoid the "big transformation trap" and instead deploy solutions that deliver value fast while building long-term capability.

CHAPTER 9

Co-Innovation Becomes the Preferred Adoption Model

The insurers moving fastest are co-building AI workflows with trusted partners. They combine:

- AWS as the operational platform
- NVIDIA as the technical accelerator
- Firemind as the insurance-focused delivery partner

This creates a loop of speed, safety, and learning. Insurers do not start from scratch. They build on patterns that are already working elsewhere, and they gain confidence through collaboration.

Co-innovation reduces risk.
It accelerates time-to-value.
It ensures internal teams grow capabilities instead of depending on vendors.
It also aligns technology with regulatory needs from the start.

CHAPTER 10

The New Operating Standard

Insurance operations will soon be:

- AI-ready
- Workflow-driven
- Governed and explainable
- Platform-supported through AWS Underwriting Workbench
- Technically accelerated by NVIDIA
- Implemented and operationalised through experts like Firemind
- Continuously improved through co-innovation

This combination will set the performance baseline for the next era of insurance operations.
Leaders who adopt this ecosystem early will move faster, adapt better, and deliver stronger outcomes across underwriting, claims, and assistance.

Conclusion

Insurance is entering a new operational era.
AI makes it possible to remove friction without replacing systems or expanding teams. It amplifies expertise, speeds up decision-making, improves consistency, and reduces cost.

The path forward is practical and proven.

- Start with one high-value workflow.
- Demonstrate value quickly.
- Scale the solution across the organisation.
- Build internal confidence and capability along the way.

The opportunity is significant and the risk is manageable.
The teams who move early will set the pace for a sharper and more competitive insurance industry.

THE FIREMIND WORKFLOW

How We Deliver AI That Works.

Insurance AI succeeds when it follows a clear, repeatable process.
Firemind uses a streamlined workflow designed for document-heavy, regulated industries like insurance, financial services, and healthcare.

Here is the simplified version.

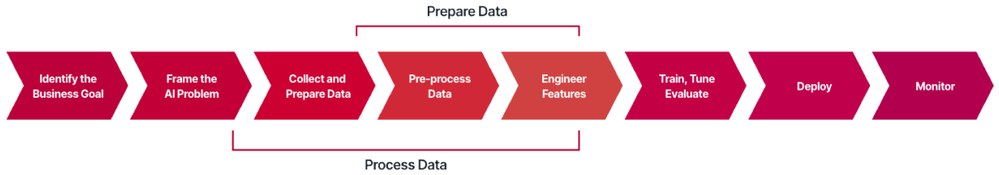

1. **Identify the Business Goal**
   Start with the outcome, not the model. We define one measurable objective, such as reducing underwriting prep time or accelerating claims triage. This keeps the project focused and ensures visible value.

2. **Frame the AI Problem**
   We translate the business goal into a clear AI task. This means deciding whether the workflow needs summarisation, extraction, classification, ranking, or generation. Clear scoping prevents overbuilding and reduces risk.

3. **Collect and Prepare Data**
   We gather the real operational documents that drive your workflow. This includes submissions, policies, claims files, medical notes, and emails. We clean and standardise everything so the AI can understand it.

4. **Train and Evaluate**
   We prototype several AI approaches and test them against real cases. Your experts validate accuracy, clarity, and safety. Only models that meet your operational standards move forward.

5. **Deploy into Operations**
   We integrate the workflow into the tools your teams already use, including legacy systems and the AWS Underwriting Workbench. Rollout is phased and monitored to ensure safe adoption.

6. **Monitor and Improve**
   AI is never “set and forget.” We track accuracy, detect drift, and apply updates based on new data or regulatory requirements. This keeps your AI reliable, consistent, and audit-ready.

Why this Approach Works.

This workflow delivers AI that is:

- Fast to implement
- Safe to adopt
- Aligned to operational reality
- Proven to scale horizontally across workflows

It is the foundation behind every Firemind insurance project: clear goals, practical delivery, and AI that makes a measurable difference from day one.

In partnership with

NVIDIA | aws