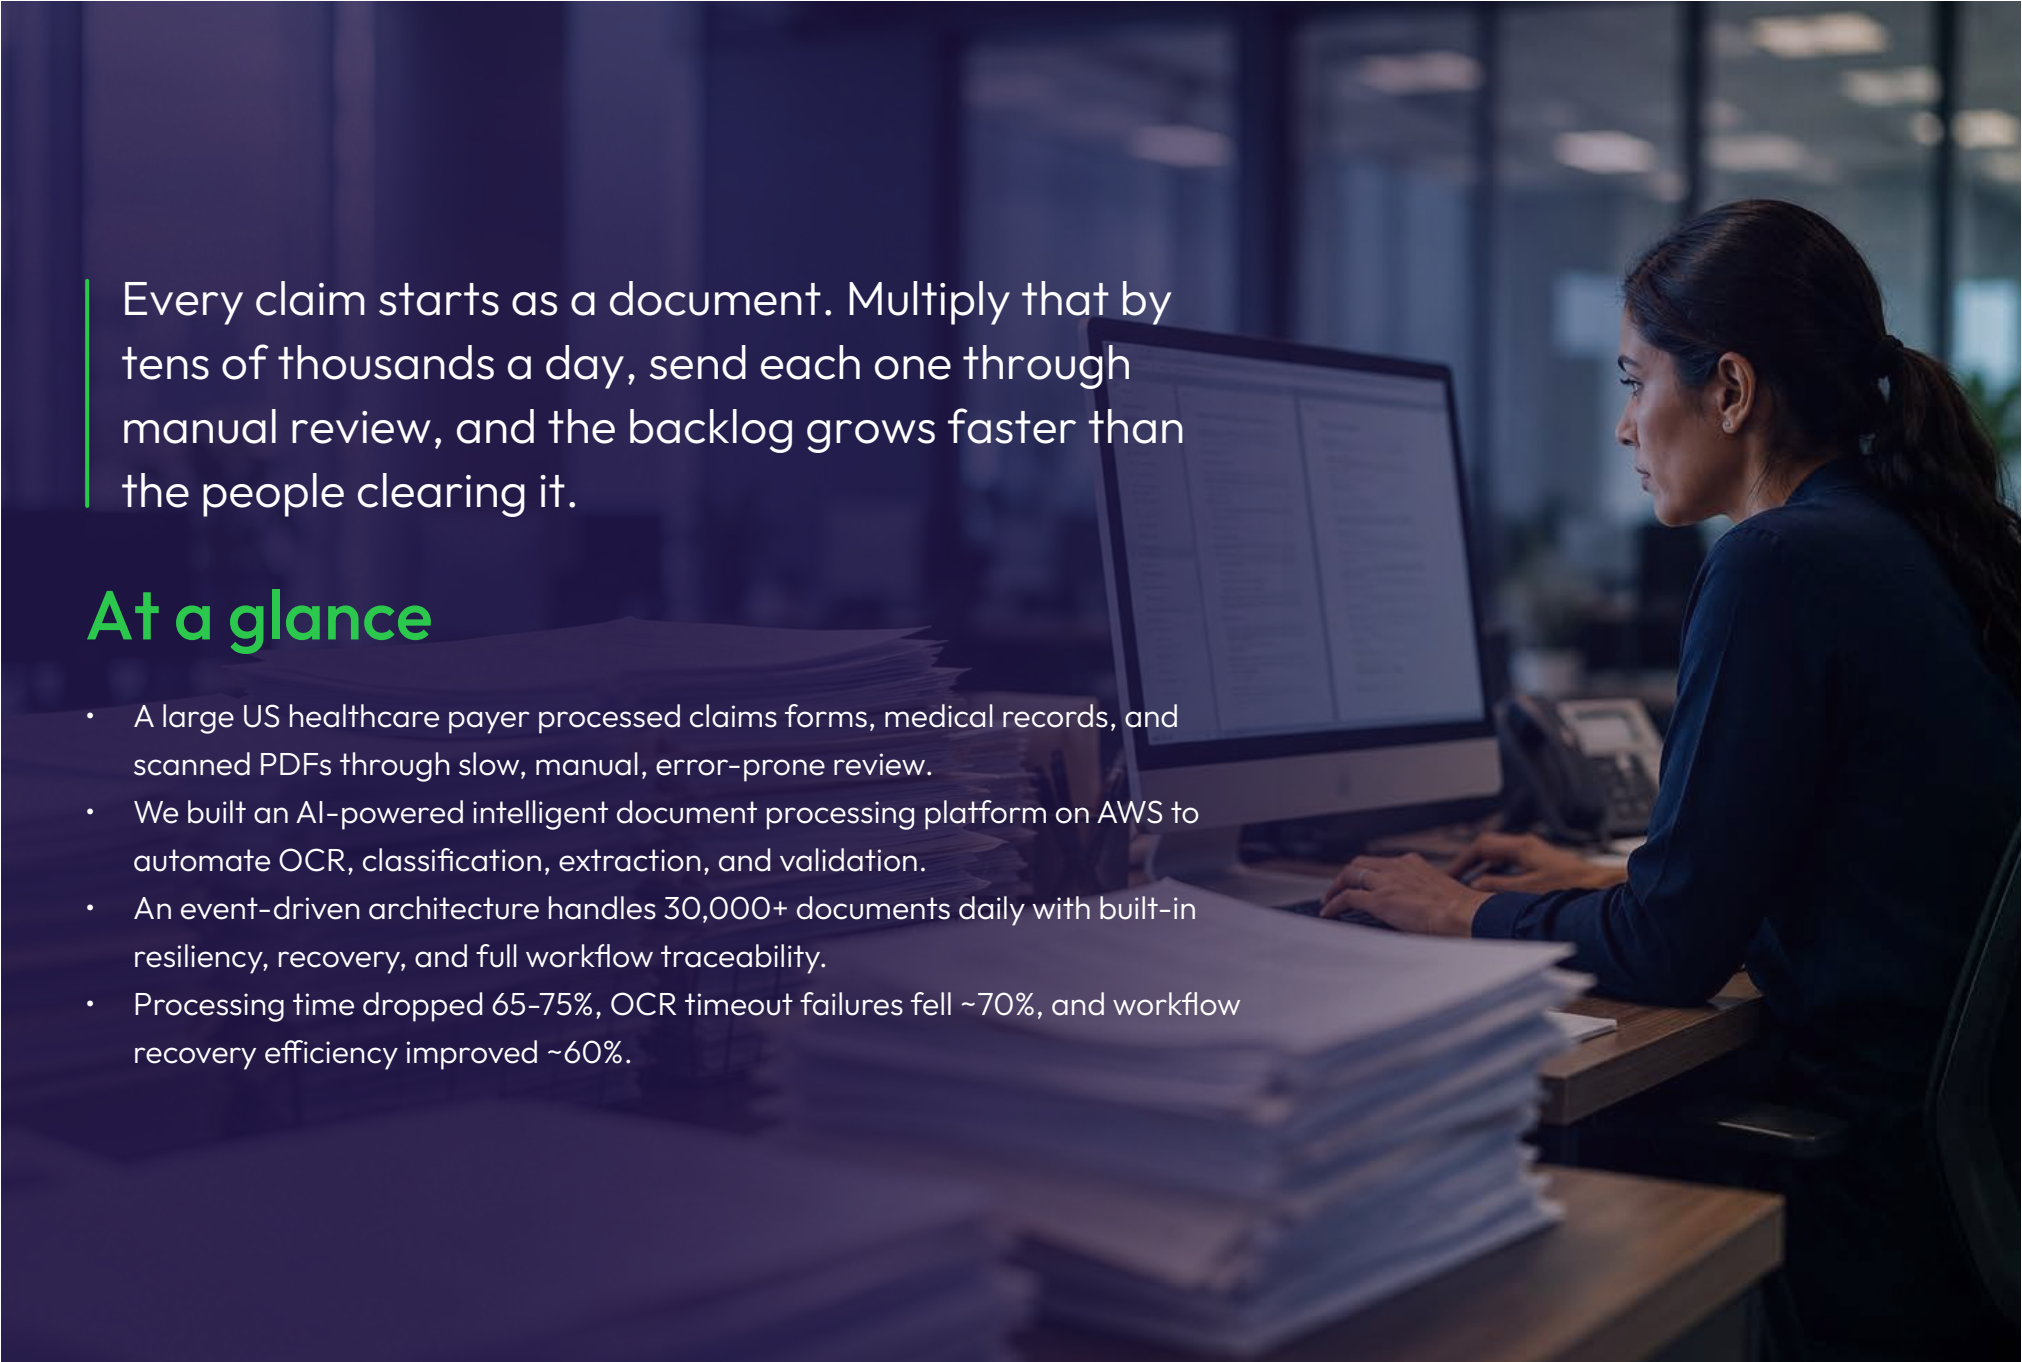




Healthcare Payer Reduces Review Time by 65–75% with AI-Powered Claims Automation on AWS





Every claim starts as a document. Multiply that by tens of thousands a day, send each one through manual review, and the backlog grows faster than the people clearing it.

At a glance

- A large US healthcare payer processed claims forms, medical records, and scanned PDFs through slow, manual, error-prone review.
- We built an AI-powered intelligent document processing platform on AWS to automate OCR, classification, extraction, and validation.
- An event-driven architecture handles 30,000+ documents daily with built-in resiliency, recovery, and full workflow traceability.
- Processing time dropped 65-75%, OCR timeout failures fell ~70%, and workflow recovery efficiency improved ~60%.

Buried under the paper trail

A large US-based healthcare payer and benefits management organization operating at enterprise scale, processes high volumes of claims forms, medical records, scanned PDFs, TIFFs, and supporting correspondence across multiple lines of business – all in a tightly regulated environment where accuracy and traceability are not optional.

For years, that work leaned heavily on people. Documents arrived unstructured, with limited metadata, and moved through manual review before they could be routed, matched, and adjudicated. It worked, until it stopped scaling. As volumes grew, turnaround times slowed, information capture grew inconsistent, and downstream claim workflows backed up behind a manual bottleneck. Every added document raised the cost of the next one – more effort, more rework, more exceptions to handle, and less visibility across the claims lifecycle.

The real problem was that the business had no reliable way to turn documents into structured, workflow-ready data.

An assembly line for unstructured documents

We built an AI-powered intelligent document processing platform on AWS, designed to do in software what had been done by hand. Each document now moves through an automated pipeline: optical character recognition, classification, entity extraction, summarization, validation, and downstream claim-matching - structured stages that turn a scanned page into enriched, workflow-ready data.

Under the surface, the platform runs as containerized FastAPI microservices on Amazon EKS, with Amazon SQS orchestrating the flow and handling retries when a step stalls. Workflow state is tracked in Amazon RDS for PostgreSQL, OCR artifacts and logs live in Amazon S3, and the AI-enriched outputs land in MongoDB. Redis-based checkpointing lets the pipeline resume from where it left off rather than starting over - a small design choice with an outsized impact on resiliency.

We delivered the platform end-to-end, from design and AI orchestration through deployment and operationalization. Releases moved through development, integration, staging, and production using Git, Tekton CI pipelines, Docker, Amazon ECR, Helm, and Amazon EKS, backed by testing, monitoring, IAM-based access controls, federated authentication, encryption, and resilience mechanisms built for a compliance-sensitive environment.

From bottleneck to throughput

The platform now processes more than 30,000 documents a day – the kind of volume that used to define the backlog, handled as routine. Document processing time fell by 65 to 75%. OCR timeout-related failures, once a steady source of stalls and rework, dropped by roughly 70%. And because the pipeline recovers rather than restarts, workflow recovery efficiency improved by around 60%.

The gains compound beyond the headline numbers. Manual review effort is down. Claims arrive downstream cleaner and more complete, ready to route, match, and adjudicate. And in an environment where regulators ask hard questions, every document now carries the traceability and operational visibility the old manual model could never guarantee.

The numbers behind the shift

30,000+

documents processed daily

65–75%

reduction in document processing time

~70%

fewer OCR timeout-related failures

~60%

improvement in workflow recovery efficiency



ABOUT BRILLIO

Brillio is The Enterprise AI Accelerator helping Fortune 1000 companies move from AI ambition to scaled impact, faster. Powered by our AI accelerator platform – Agentic Data and Application Management (ADAM), Brillio is one of the fastest-growing digital technology service providers, delivering transformation across five core workstreams: business-led transformation, customer experience transformation, AI and data engineering, digital engineering, and infrastructure engineering.

With 14 delivery locations across North America, Europe, and Asia and a team of over 6,000 customer-obsessed professionals, Brillio combines deep industry expertise, modern engineering, and accelerators to deliver measurable outcomes. Headquartered in Dallas, Texas, Brillio serves clients globally with a commitment to speed, scale, and measurable impact.



<https://www.brillio.com/>

Contact Us: info@brillio.com