


GENAI DOCUMENT INTELLIGENCE IMPLEMENTATION

Every document read. Every field on **record**.

Most teams still retype invoices, claims, and contracts by hand, and the cost grows with every document. We build one auditable workflow on Amazon Bedrock and Amazon Textract that turns them into structured data.

HOW IT WORKS

One workflow, inside your AWS account



Inside your security boundary

The pipeline is deployed into your AWS account. Documents never leave it.



Human review built in

Uncertain fields go to a person instead of failing silently, so accuracy improves over time.



Every decision logged

Each extraction is recorded with its confidence score, and the infrastructure is code your team owns.

HOW WE ENGAGE

Scoped to your document set

- Discovery and design**
Target architecture, measured baseline, and a costed plan.
- Pilot implementation**
One document class end to end, in a non production account.
- Production rollout**
Full document set, monitoring, runbooks, and handover.
- Managed run and optimize**
Ongoing evaluation, accuracy tuning, and monthly reporting.

Scope and timing are set with each customer.

CASE STUDY

A regional property and casualty insurer automated claims intake

62%

drop in average claim intake time¹

3x

claim volume handled per reviewer, per shift²

1. Insurer operations report, Q2 2026 (e2e test data). 2. Insurer staffing model, June 2026 (e2e test data).

Talk with an expert at sales@avahi.ai

Every engagement has a named delivery lead. After handover, a 30 day support window covers questions, defect fixes, and runbook clarifications.

AWS Premier Tier Services Partner
avahi.ai | © 2026 Avahi