

Took engineering teams from one-off AI experiments to a governed lifecycle on Amazon Bedrock — every feature evaluated and audit-logged before production

Aviva — Multi-week programme

We delivered a hands-on AI development lifecycle programme for a major insurer, moving teams from ad-hoc experimentation to a single governed path to production on Amazon Bedrock — cutting time-to-first-governed-deployment by an estimated 50% while embedding the evaluation, model-versioning and audit logging a regulated insurer requires.

THE CHALLENGE

A large insurer had strong AI ambitions but no shared, repeatable way to build generative-AI features safely. Teams were experimenting in isolation, with no consistent evaluation, no model-versioning and no audit trail — meaning nothing could be trusted into production in a regulated insurance setting. Generic vendor training would have taught patterns, but not the governance, model-risk and traceability practices the insurer is held to.

WHAT WE BUILT

- Designed an AI development lifecycle programme covering the full path from prototype to governed production on Amazon Bedrock
- Ran hands-on sessions on prompt engineering, RAG patterns, evaluation harnesses and guardrails, grounded in real insurer use cases
- Embedded governance from the start: model and version control, evaluation against golden datasets, human-oversight checkpoints for sensitive decisions, and audit logging of AI behaviour
- Left teams with a reusable internal playbook and reference patterns so every new AI feature follows the same governed, auditable lifecycle

OUTCOMES

- Cut time-to-first-governed-deployment by an estimated 50%, with every AI feature now evaluated against golden datasets and audit-logged before it ships
- Engineering teams equipped with a single repeatable, governed lifecycle on Bedrock rather than isolated experiments
- Shared vocabulary and governance patterns across previously siloed teams