AI ARCHITECTURE + GOVERNANCE

Know where your agents can act. Know what to improve.

An evidence-based technical assessment for organizations developing, deploying, or scaling AI agents on AWS.

WHAT YOU RECEIVE

A technical report your engineering team can act on

Architecture observations | prioritized findings | implementation alternatives | 30/60/90-day roadmap

THE ASSESSMENT

The questions we help answer

- What can agents access, invoke, or change?
- Where should authorization and human review occur?
- How are tool actions and decisions traced?
- What happens when actions fail or are retried?
- Which controls belong in AWS, custom code, or dedicated tooling?

A clear five-stage engagement

- 01 Intake | Secure questionnaire and documents
- 02 Discovery | 60-minute technical call
- 03 Assessment | Evidence-based review
- 04 Report | Written findings and roadmap
- 05 Advisory | 60-minute findings call

STANDARD ENGAGEMENT

1	1	Up to 2	2 x 60 min
defined AI agent system	primary AWS environment	agent workflows	advisory calls

THE DELIVERABLE

Assessment: Security, permissions, governance and operational controls. Options: AWS-native, custom and third-party paths with tradeoffs. Roadmap: Priority actions for the next 30, 60 and 90 days.

Clearly scoped. No mandatory software purchase.

No penetration testing, compliance certification, full source-code audit, implementation, or managed operations. Scope and fees are agreed through an AWS Marketplace private offer.