



Integrating Claude on AWS & Maximizing ROI

Unlocking frontier AI for your business

A practical guide for enterprise technology and data leaders: what changed, why Claude, and why a 90-day roadmap from pilot to production.

What you'll find inside

01	Executive summary The case in one page	06	The first 90 days A pragmatic adoption roadmap
02	The landscape has changed From chatbots to agentic systems	07	Governance & responsible AI What your CISO reads first
03	Why Claude for the enterprise Reasoning, safety, flexibility	08	Why a services partner What AICG brings
04	Deployment architectures Anthropic API vs. Amazon Bedrock	09	Getting started The first conversation
05	Five highest-value patterns Where enterprise value concentrates		

Enterprise AI has crossed a threshold

The question is no longer *whether* large language models belong in the enterprise stack — it is *which* models, deployed *where*, governed *how*, and integrated with *whose* data.

Model selection is the easy part. Organizations seeing real returns connect AI to governed data, start with narrow measurable workflows, and treat evaluation and governance as engineering disciplines from day one.

KEY TAKEAWAYS

Agentic AI is where value concentrates

Models that plan, use tools, and execute demand stronger data foundations than chat ever did.

Bedrock keeps AI in your boundary

Adopt frontier AI without inference leaving your cloud account boundary.

Architect for MCP now

The Model Context Protocol is the emerging standard interface between AI and enterprise systems.

Budget for the demo-to-production gap

The gap is evaluation, observability, and governance. Fund it from the start.

From chatbots to systems that do the work

From answers to actions

Frontier models plan, call tools, query databases, and execute multi-step workflows with human checkpoints. Delegate processes, not just prompts: pipeline triage, reconciliation, board-ready reporting.

From generic to enterprise context

RAG, long context windows, and standardized protocols mean the model works with your contracts, your sales data, your documentation — with citations and audit trails.

From experiments to accountability

Boards and regulators now ask: what data does the model see, who reviews outputs, how is it evaluated? Mature answers ship. The rest stall in pilot purgatory.

 **“The differentiator is no longer the model — it is the quality and accessibility of the data behind it.”**

Competitive-to-leading on all four axes

Capability, safety, deployment flexibility, and ecosystem trajectory.

Frontier reasoning & coding

At or near the top of independent evaluations for complex reasoning, long-document analysis, and software engineering — including Claude Code. Fewer hallucinations on dense material.

Long context, real documents

Entire contracts, codebases, filings, or policy manuals analyzed in a single pass — no brittle document-chunking gymnastics.

Safety-first by design

Constitutional training, strong refusal behavior, calibrated uncertainty. For finance, healthcare, and insurance, this posture materially lowers deployment risk.


Deployment flexibility

Anthropic API directly, or native on Amazon Bedrock — the same IAM, networking, and billing you already govern. Vertex AI for multi-cloud estates.



Ecosystem trajectory: MCP. Anthropic created the Model Context Protocol, now adopted broadly. Build MCP servers over your systems once — any compliant model can use them. Architecture for the next five years, not the next release cycle.

Two paths to the same models



A governance and operations decision, not a capability one.
Many enterprises run both.

Where enterprise value concentrates

01

Document intelligence at scale

Contracts, invoices, claims, filings — read, extracted, classified, summarized with citations. Enter via one high-volume workflow with measurable cycle time.

02

Analytics augmentation

Claude atop the governed warehouse — via MCP or semantic layers — answering business questions, drafting recurring reports. Prerequisite: trustworthy data.

03

Agentic workflow automation

Pipeline triage, reconciliation, vendor onboarding, ticket resolution. The agent proposes, a human approves, the system logs everything.

04

Engineering acceleration

Claude Code across generation, review, migration, documentation. Often the fastest measurable ROI — engineering output is already instrumented.

05

Customer experience

Support deflection, agent-assist, personalized communication — grounded in your knowledge base through RAG, with guardrails and escalation designed in.



“Most stalled AI initiatives don't have a model problem. They have a data problem wearing an AI costume.”

A pragmatic adoption roadmap

 DAYS 0–30

Assess & aim

- AI readiness assessment: data estate, security posture, team skills
- Use-case inventory scored on value, feasibility, data readiness — pick one or two, not ten
- Deployment decision (API / Bedrock / hybrid) with security at the table
- Success metrics defined *before* building

 DAYS 31–60

Build the thin slice

- Production-shaped pilot: real data, real auth, real logging — scoped narrow
- Integration layer: MCP servers or RAG pipeline over governed sources
- Evaluation harness built in parallel — golden datasets, regression tests
- Human-in-the-loop checkpoints designed in, not bolted on

 DAYS 61–90

Prove & scale

- Measured rollout to a defined cohort; compare against baseline
- Governance review: audit logs, access controls, incident playbook
- Cost model validated: caching, batching, model-tier routing
- Next three use cases sequenced by shared infrastructure

Zero to a measured, governed production pilot — if scope stays narrow and metrics come first.



The output of 90 days is not a demo — it is a governed, measured production pilot, an evaluation harness, and a sequenced roadmap your leadership can fund with confidence.

The section your CISO reads first

None of this is exotic — it is the same engineering discipline you already apply to yourPII, extended to model calls.

 **Logging is what turns “the AI did it” into an auditable business process.**

THE ENTERPRISE AI GOVERNANCE CHECKLIST - ABRIDGED

- | | |
|--|---|
| ☐ Data boundary clarity
What data reaches the model, under what agreement, with what retention. | ☐ Access & identity
Role-based access via MCP permissions — no god-mode service accounts. |
| ☐ Evaluation as a discipline
Golden datasets, automated regression, human review sampling. | ☐ Observability & audit
Every call logged: inputs, outputs, tools, approver identity. |
| ☐ Usage policy alignment
Map acceptable use to Anthropic's policy, SOC 2, HIPAA, model-risk guidance. | ☐ Human accountability
Agents propose; accountable humans approve. Design overrides before launch. |

You could do all of this yourself, but...

The question is whether you should spend two quarters learning lessons a partner has already learned for you. AICG brings three things generic AI consultancies or consultancies try to say they have an AI practice don't have:

Data engineering depth

We build the warehouses, lakehouses, pipelines, and platforms — including [DLH.io](#) and dbdeux dbt — that AI stands on. We fix root causes, not prompt around them.

Production engineering discipline

Versioned, tested, observable systems your team can own: CI/CD, evaluation harnesses, cost governance, documentation, training.

Right-sized engagement

Strategy only, full build, or embedded augmentation — small engagements are viable and palatable, large ones require efficiency only a knowledgeable partner can generate - that's AICG.

09 • GETTING STARTED

Start with the two or three workflows where AI could move a number your leadership already tracks

That's where every successful program begins. AICG's Enterprise Claude AI Integration Services are available through AWS Marketplace via private offer.

sales@aicg.com aicg.com