MONOCLE TECH

The AI opportunity is real.
The gap to value is the problem.

2026

TECH

Most organisations are using AI. Few are benefiting from it.

• Scale Problem

<10%

of enterprises are successfully scaling AI agents beyond pilots.

The leap from proof-of-concept to production remains the single hardest step in the enterprise AI journey.

Source · McKinsey QuantumBlack, Apr 2026

• Measurement Problem

≠

Adoption is not impact.

Organisations measure tool usage, not business outcomes. The result: impressive dashboards, disappointing results.

Pattern observed across 40+ Monocle engagements

• Structural Problem

−19%

slower with AI tools - when the surrounding process doesn't change.

Bolting AI onto unreformed processes makes experienced teams slower, not faster. The constraint is structural, not talent.

Source · METR developer-productivity study, 2025

TECH

The four tensions every AI leader needs to resolve.

Adoption vs Effort

Reclaimed time becomes slack, not output.

Employees use AI to finish faster - and stop there. Adoption dashboards look healthy while higher-value work goes unstarted.

Adoption vs Value Realisation

Deploying tools is not the same as changing outcomes.

Most organisations can count seats and queries. Few can show the P&L line that moved because of them.

Augmentation vs Automation

Automating junior roles removes the talent pipeline.

Automate chaos and you get chaos at scale. Automate the apprenticeship and you lose tomorrow’s senior bench.

Machine Productivity vs Human Judgment

AI excels at structured cognition. Judgment is the differentiator.

Asking better questions, evaluating output, deciding what matters - these remain decisively human, and decisively scarce.

Where the Delivery Complexity Actually Sits

TECH

The model is the smallest part of the system.

A credit-analyst co-pilot needs an LLM - and eight other things Monocle already knows how to build.

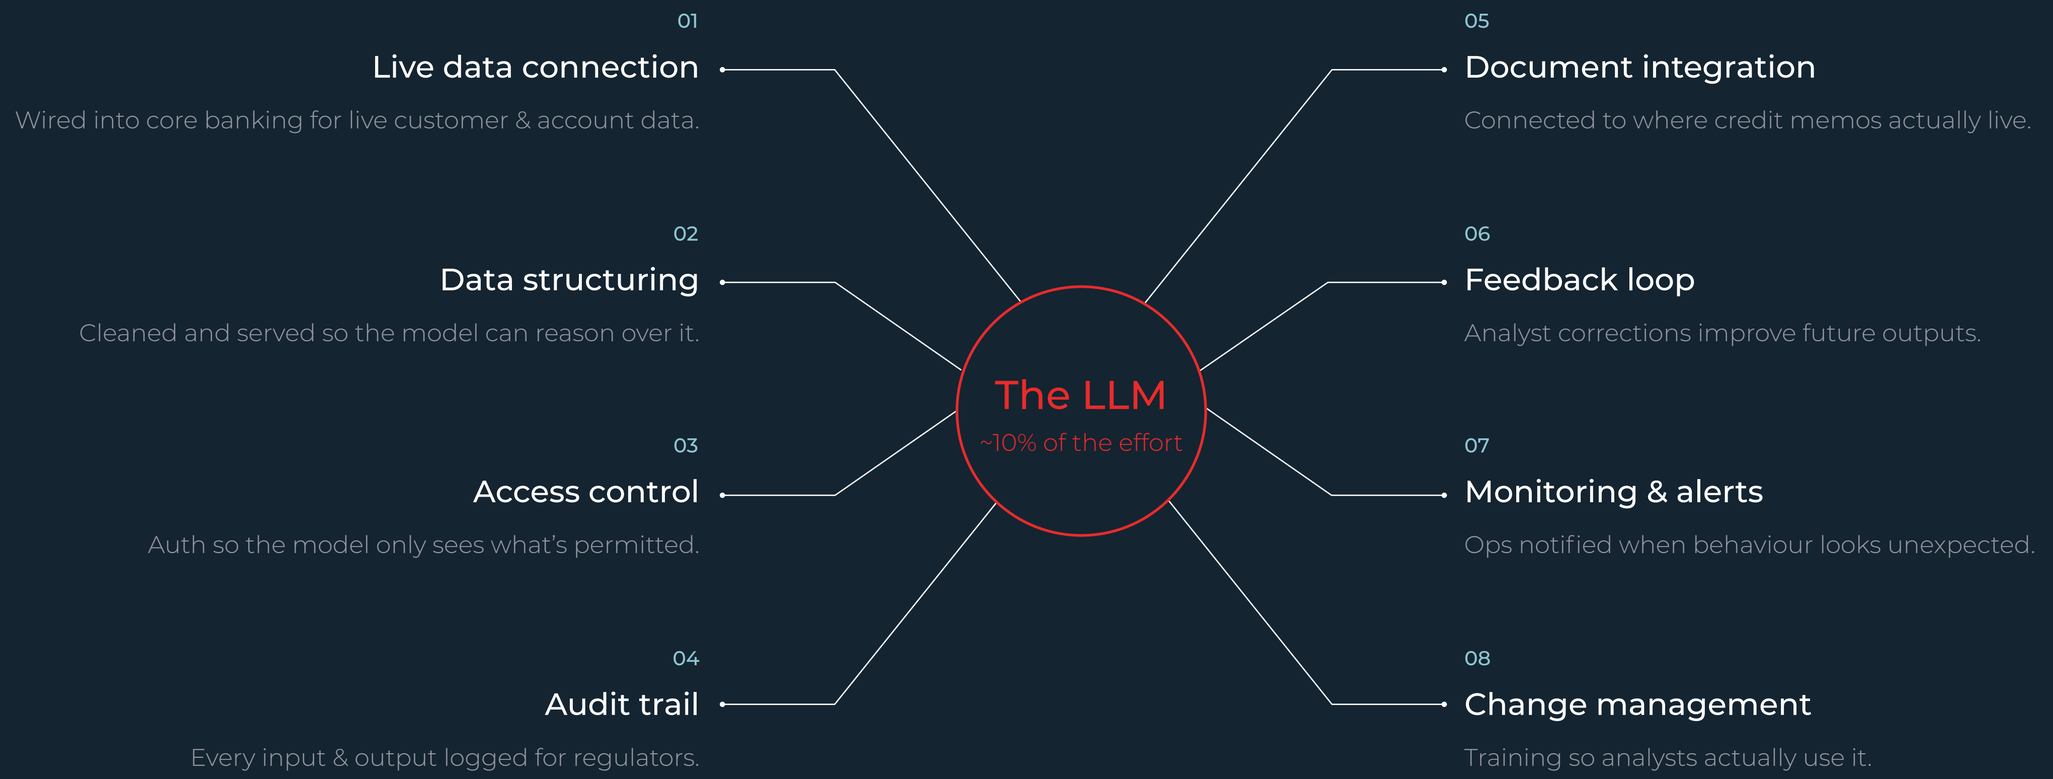

From Pilot to Production

TECH

Agentic AI is in production at leading banks. The gap is architectural.

Case · European Bank

Credit application agent

An agentic workflow handles routine credit assessment steps end-to-end, freeing relationship managers for the work clients actually pay for.

20%+ RM time reclaimed

Case · Global Financial Services

Agent factory · 100+ developers

A reusable agent platform supports COBOL and Java modernisation across the bank - every team building on the same governed substrate.

50% faster modernisation

Case · Global Bank

Overnight engineering sprints

AI agents run on software development tasks through the night, delivering continuous progress without continuous headcount.

24/7 continuous delivery

Enterprise AI Value Chain

Agentic layer - operates across the chain

01 Strategy	02 Architecture	03 Build	04 Deploy	05 Operate

Agentic Overlay — Agents reach across the value chain - planning, building, deploying and observing - when the architecture lets them.

The question is no longer whether agentic AI works at scale. It is whether your architecture is ready to support it.

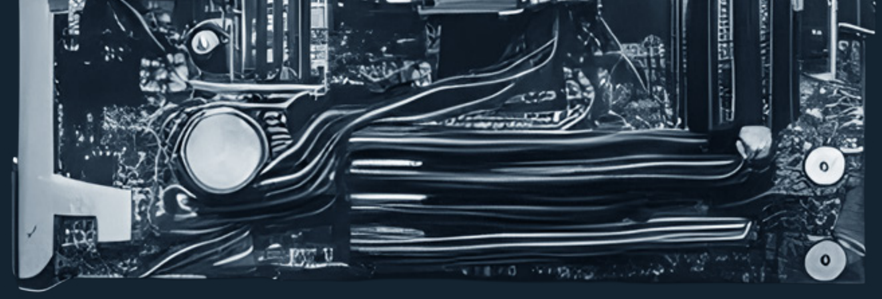

TECH

What we offer.

An integrated set of solutions across the full data and AI value chain.

01.

+ Data Engineering & Modern Data Platforms

Lakehouse architectures on open table formats (Iceberg, Delta, Hudi) supporting batch and streaming workloads.

02.

+ Software Engineering & Solution Delivery

Full-stack engineering teams building secure, scalable applications and APIs that operationalise data and AI inside the bank.

03.

+ Machine Learning & GenAI

From predictive modelling and MLOps through to enterprise-grade GenAI capabilities deployed with strict privacy and security controls.

04.

+ Cloud Engineering, DevOps & FinOps

Cloud-native platforms across Azure, AWS and GCP, engineered for scale, automation, and cost-effective operation.

05.

+ Data & Solution Architecture

The connective tissue across every capability above. Target-state blueprints, reference architectures and governance patterns that align business intent with regulated, cloud-native delivery.

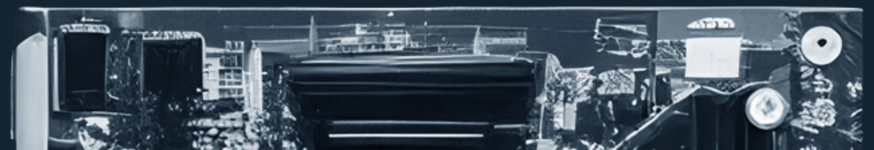

TECH

From AI ambition to production.

Strategy → Delivery → Governance
One engagement model, scoped to where you are.

01
AI Strategy & Roadmap
Building credible, governed AI programmes with a clear path from ambition to production.

02
AI Use Case Discovery & Prioritisation
Systematic identification and scoring of AI opportunities across the business - structured, repeatable, high value per day.

03
AI Solution Design & Build
Engineering AI-powered systems grounded in four core architecture patterns: RAG, co-pilot, agentic workflow, and classification.

04
AI Programme Delivery & Governance
Managing AI delivery at scale with appropriate oversight, audit trails, and regulatory alignment.

05
AI Capability Building
Upskilling client teams to sustain AI adoption, directly informed by Monocle's own Role Champion model and internal experience.

06
AI Governance & Risk Frameworks
Compliance-driven governance, ethics, and risk frameworks for regulated clients - aligned to FCA, PRA, and EU AI Act requirements.

07
AI Enterprise Enablement
Making AI work inside the client's environment - connecting AI systems into existing banking infrastructure.

- AI Connectivity & Data Enablement
- AI System Integration & Architecture
- AI Operationalisation & Change

TECH

AI is not the answer to everything. Choosing the right intelligence is.

We match the right solution - human, rules, ML, or GenAI - to each problem.

Human judgment

Ambiguity & accountability

Humans own high-stakes, ethical, and ambiguous decisions.

- Credit & underwriting exceptions
- AML/SAR investigation decisions
- M&A & capital allocation calls

Rules & code

Known, stable logic

Deterministic, explicit “if X, then Y” logic for predictable outcomes

- Payment & settlement processing
- Regulatory reporting calculations
- KYC data validation

Machine Learning

Patterns in structured data

Statistical prediction where patterns exist but rules are too complex to hand-write.

- Card & payments fraud detection
- Credit risk & churn scoring
- Liquidity & demand forecasting

Generative ai

Reasoning over unstructured input

Flexible interpretation and language tasks where some error is tolerable.

- RAG over policy & compliance docs
- Research & variance commentary
- Credit application co-pilots

Orchestration

Agentic AI: bringing all four together

The strongest solutions aren’t AI-first; they’re hybrid systems orchestrating multiple forms of intelligence.

- Agent as orchestrator, not a single model
- Deterministic tools for anything that must be exact
- Human-in-the-loop for accountability

TECH

We bridge the divide between business and tech.

AI-Native Firms

Deep on the model.
Thin on the enterprise.

Strength

Understand models, agents, and LLM architecture deeply.

Gap

Often lack the integration experience to connect AI to legacy-heavy, regulated banking infrastructure.

The Gap

Monocle

Understands the AI layer well enough to make the right integration decisions - and the enterprise layer well enough to execute them in a bank, without Big 4 overhead.

Large SIS & Consultancies

Deep on the enterprise.
Thin on the model.

Strength

Enterprise integration experience across complex banking platforms.

Gap

Often slow, expensive, and still building AI depth fast enough to be credible.

TECH

Four ways Monocle Tech accelerates your AI journey.

From strategy through to production.

01 · AI Service Offering

From strategy to production.

Strategy & roadmap, use case discovery, solution design & build, programme delivery & governance, capability building, risk frameworks, and enterprise enablement - one partner across the full journey.

02 · AI-Enabled Delivery

AI-augmented teams that ship, faster and leaner.

We deliver consulting and technology projects using AI-augmented teams - greater throughput than traditional models, with the engineering rigour regulated environments require.

03 · AI Accelerators

Pre-built frameworks to compress pilot-to-production.

Emerging methodologies, assessment approaches, and reusable assets informed by our hands-on AI delivery experience, helping clients accelerate readiness assessments, identify value opportunities, and avoid starting from a blank slate.

04 · AI Capability Building

Competency that stays inside your organisation.

Role-specific playbooks, governance practices and hands-on enablement so AI maturity outlasts our engagement - and your teams own what comes next.

TECH

Foundations for making AI production-ready

01 / Data Readiness

The most underestimated prerequisite

AI systems amplify data quality problems rather than hiding them. Most banks have a data platform - few were designed with AI consumption in mind.

- Accessibility, quality & freshness
- Lineage & provenance
- PII controls & masking

02 / Integration Architecture

Where pilots stall after launch

AI must connect to where work happens - core banking, CRM, case management. This builds directly on Monocle's existing integration capability.

- API quality & event streaming
- Bidirectional integration & write-back
- Latency & governance

03 / Security & Access Control

A new class of attack surface

Agentic AI is a security boundary most frameworks weren't built for. Monocle frames this in the language of enterprise risk, not just controls.

- Least-privilege & identity for AI
- Prompt injection & exfiltration controls
- Third-party model risk

04 / Governance & Auditability

Monocle's most differentiated domain

SR 11-7, PS 7/23 and the EU AI Act make explainability a current, not future, obligation - the governance wrapper regulators will accept.

- Audit trails & model versioning
- Explainability & human override
- Bias monitoring & model risk management

05 / Observability & Resilience

AI degrades differently to software

Models fail silently and drift gradually. Standard IT monitoring isn't sufficient.

- Drift & hallucination detection
- Cost monitoring & alerting
- Business continuity fallback design

06 / Operating Model & Change

Works in UAT, fails in production

Ownership, human-in-the-loop design and AI literacy, made more urgent by the AI context.

- Ownership & accountability
- Staff readiness & feedback loops
- Vendor & model dependency

Making AI Work in your Environment

TECH

Three service lines. One integration capability.

Service Line 01 · Foundational Layer

AI Connectivity & Data Enablement

- Data readiness assessment for AI
- Data pipeline design & build
- API & integration layer design
- Data governance alignment

Service Line 02 · Middle Layer

AI System Integration & Architecture

- AI architecture design - model, retrieval, memory, tools, orchestration
- Identity & access management integration
- Workflow & case management integration
- Audit & logging architecture

Service Line 03 · Sustaining Layer

AI Operationalisation & Change

- Operating model design for AI systems
- Human-in-the-loop process design
- Model monitoring & observability
- Change management & adoption

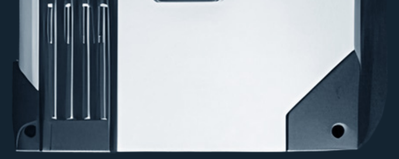

TECH

The right partner for regulated, complex environments.

Banking & enterprise DNA

We understand the constraints of regulated, legacy-heavy environments. We do not optimise for greenfield.

15+ Years Inside Tier-1 Banks

Delivery, not just advisory

We build alongside you. Our teams deliver working systems in production, not slide decks describing them.

Outcomes Shipped, Not Authored

Reusable assets, not billable hours

Our accelerators compress timelines. You benefit from frameworks built across multiple engagements without paying for them each time.

Accelerator Library Across 40+ Projects

Governance built in

AI in regulated industries requires accountability. We design for auditability, explainability and responsible-AI from day one.

Audit-Ready By Design

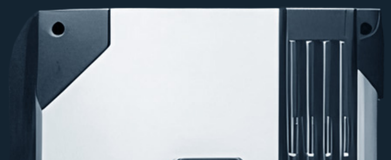

Insights → design → build → embed.

A disciplined four-phase lifecycle, anchored in business value.

Phase 1

Insights & Readiness

Identify high-impact business opportunities and run rigorous data maturity and risk assessments to ensure the underlying data is accurate, compliant, and AI-ready.

→

Phase 2

Design & Proof of Concept or MVP

Agile consultants execute rapid prototyping and architectural design - validating technical feasibility and business value through low-cost POCs or MVPs.

→

Phase 3

Build & Productionise

Transition validated concepts into scalable enterprise solutions using best-practice DevOps and MLOps, with strict security guardrails and privacy controls.

→

Phase 4

Embedment & Change

Drive sustainable adoption through comprehensive change management and staff upskilling - embedding solutions in the day-to-day operating rhythm.

Technology-agnostic · Outcome-obsessed

POC → Production → Post-production Support

TECH

GenAI deployments at leading banks.

Use Case 01 · Finance

Variance commentary at scale

An LLM-powered insights engine generating customised variance commentary from General Ledger data - using tailored prompts to explain monthly movements for a retail, corporate and investment bank's finance control team.

~30% less time

Reduction in routine monthly reporting time across hundreds of report stakeholders.

Use Case 02 · Research

GenAI research assistant

A research workflow transformation at a leading investment bank - drafting research reports, summarising market commentary, highlighting trends, and supporting analyst queries from both internal and external sources.

Analyst uplift

Faster delivery of insights; analysts redirected to higher-value client engagement.

Use Case 03 · Credit

Credit application auto-complete

An adaptive-prompt LLM solution streamlining credit applications for a retail, corporate, and investment bank - auto-completing forms from internal data and contextual cues, with robust data privacy safeguards.

Faster turnaround

Reduced manual entry and accelerated application cycle times.

Use Case 04 · SDLC

AI-native software delivery

Transformation across the software development lifecycle using Kiro, an AI-powered development environment turning requirements into specs, technical designs and production-ready code.

Faster delivery

Spec-to-code workflow compressing design, build and implementation cycles.

Engagement Pathway

TECH

You don’t need to boil the ocean. You need the right entry point.

Diagnose → Design → Deliver
A three-step path from clarity to production.

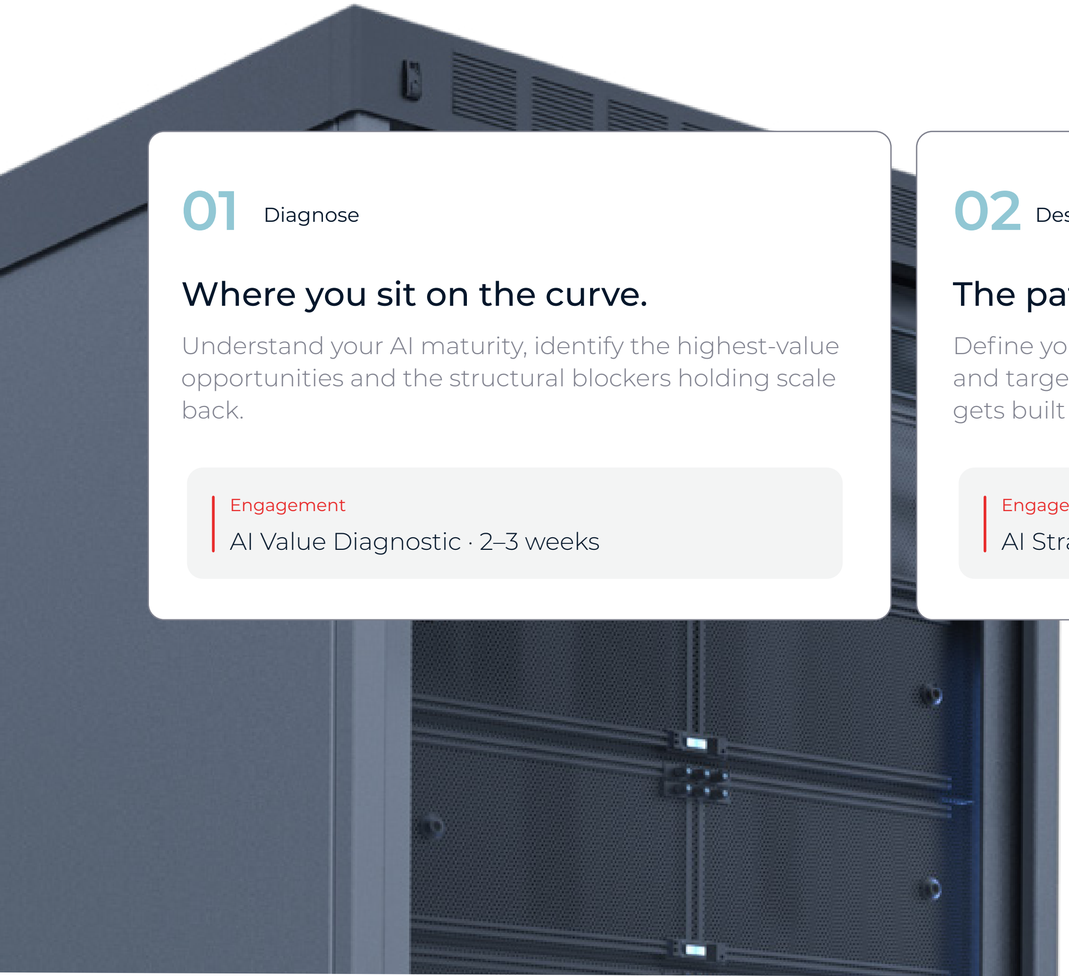

01 Diagnose

Where you sit on the curve.

Understand your AI maturity, identify the highest-value opportunities and the structural blockers holding scale back.

Engagement
AI Value Diagnostic · 2–3 weeks

02 Design

The path from pilot to production.

Define your AI use-case roadmap, governance model and target architecture. Align stakeholders on what gets built - and what doesn’t.

Engagement
AI Strategy & Roadmap · 4–6 weeks

03 Deliver

Working systems in production.

Build, test and deploy AI-enabled solutions with Monocle’s delivery teams - with accelerators, governance and capability transfer built in.

Engagement
AI Delivery Partnership · ongoing

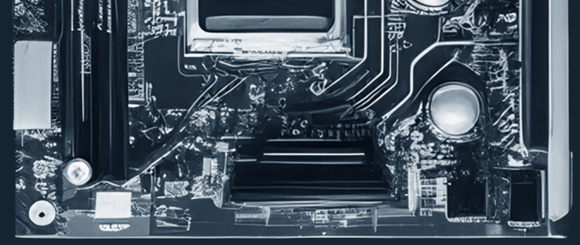

TECH

Four flexible engagement models - structured for predictability, sized for impact.

Model 1

Delivery-Based

Fixed-scope or outcomes-driven engagements with defined milestones and acceptance criteria.

- Cost predictability
- End-to-end accountability
- Platform & AI use-case delivery

Model 2

Capacity-Based

Time-and-materials consultants enabling clients to scale capacity up or down on demand.

- Staff augmentation
- Ongoing engineering support
- Specialist roles

Model 3

Dedicated Squads

Multi-disciplinary squads driving a stream of work or operating a platform capability.

- Technical leadership
- Data & ML engineers
- Cloud / BA / SME blend

Model 4

Hybrid

Combine fixed-scope components with flexible resourcing for both structure and adaptability.

- Fixed-scope build
- Ad-hoc BAU support
- Specialist AI engineers

www.monoclesolutions.com

The organisations that close the gap now will define the benchmark for the future.

AI adoption is widespread. Value realisation is rare. The window to differentiate is now - and it is narrower than the hype cycle suggests.

Johannesburg

info.za@monoclesolutions.com

+27 (0) 11 263 5600

Cape Town

info.za@monoclesolutions.com

+27 (0) 21 007 5979

London

info.uk@monoclesolutions.com

+44 (0) 2030 022 514

Amsterdam

info.nl@monoclesolutions.com

+31 (0) 2 02 256 182