

— AI READINESS ASSESSMENT · 2026

AI Readiness Assessment

An independent maturity assessment across ten domains, benchmarked to ISO/IEC 42001, Google SAIF and the NIST AI RMF. A sample report, shared to illustrate the deliverable.

2.3_{/5}

DEVELOPING

Pockets of good practice exist but they are inconsistent and fragile.

CLIENT

SampleCo XYZ – mid-market financial services, ~480 staff

DOMAINS ASSESSED

10 · People · Process · Technology · Governance

PREPARED

05 June 2026 · valid 90 days

ENGAGEMENT

AI Readiness Assessment – 100-question instrument

FRAMEWORKS


ISO/IEC 42001 · Google SAIF · NIST AI RMF

SIGNED OFF BY

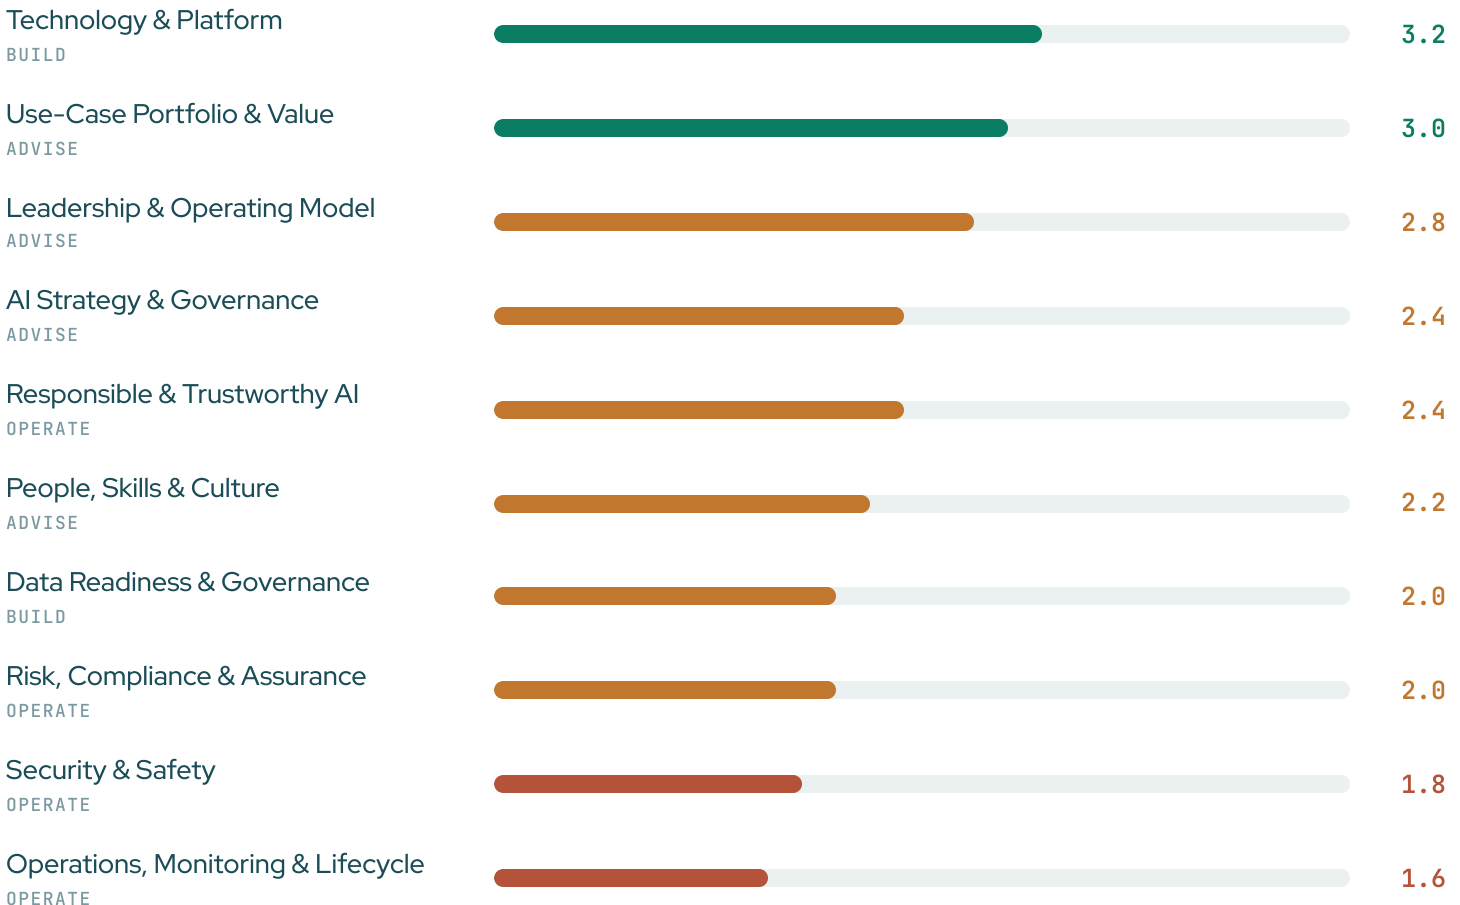
Patrick Sullivan, Managing Director

A capable business running ahead of its guardrails

SampleCo has real momentum: teams are experimenting, leadership is engaged, and the technology base is solid. But the enabling foundations – **security, risk, data governance and lifecycle operations** – have not kept pace. The result is an overall maturity of **2.3 / 5** (“Developing”): enough to create value, not yet enough to do so safely or repeatably at scale. The pattern is common and fixable. The next two quarters matter most.



Maturity by domain



▲ Where SampleCo is strongest

Technology & Platform 3.2/5

▼ Where the risk & upside sit

Operations, Monitoring & Lifecycle 1.6/5

SECTION 02

Findings

What the scores mean

BUILD

Technology & Platform

3.2 / 5

A modern, cloud-based platform and engineering team give SampleCo a genuine head-start – the rails to run AI already exist.

ADVISE

Use-Case Portfolio & Value

3.0 / 5

Teams have identified valuable use-cases and several pilots show clear promise. The ideas are sound; what's missing is the foundation to scale them safely.

ADVISE

People, Skills & Culture

2.2 / 5

Appetite is high but skills are uneven and concentrated in a few enthusiasts. Without a capability plan, delivery stays bottlenecked on individuals.

BUILD

Data Readiness & Governance

2.0 / 5

Data is fragmented across systems with inconsistent quality and access controls. This is the single biggest cap on every use-case above it.

OPERATE

Security & Safety

1.8 / 5

AI is being used faster than it is being secured. There is no SAIF-aligned control set and no AI-specific incident path – a material exposure for a financial-services firm.

OPERATE

Operations, Monitoring & Lifecycle

1.6 / 5

Once a model is live, little is watched. No drift, performance or cost monitoring means problems surface as incidents, not signals – the lowest score and the most urgent gap.

SECTION 03

Recommendations

How Millwater would help — three lanes

OPERATE

Stand up the safety & assurance baseline first

Operations (1.6), Security (1.8) and Risk (2.0) are the lowest scores and the highest exposure. We establish monitoring, an AI incident path, a SAIF-aligned control set and a lightweight risk register – so existing AI use is no longer ungoverned.

BUILD

Fix the data foundation the use-cases depend on

Data Readiness (2.0) caps everything above it. We design the data governance, access and quality layer that turns promising pilots into production-grade, repeatable systems.

From Developing (2.3) to Defined (3.0+)

Each phase lifts the lowest-scoring foundations first, because they gate everything above them. The target is a defensible **Defined** maturity inside twelve months.

- 0-3 months**
Make AI safe to use
 - AI policy & acceptable-use approved and circulated.
 - AI inventory – catalogue every model, tool and pilot in use today.
 - SAIF baseline controls on existing AI; an incident-response path.
 - Risk register stood up; a named accountable owner per use-case.

- 3-6 months**
Build the foundation
 - Data governance layer – quality, access, lineage for priority data.
 - Documented AI strategy linked to board priorities & funding gates.
 - Capability plan – role-based AI training; an internal champion network.
 - Two pilots productionised on the new foundation as proof points.

- 6-12 months**
Scale with assurance
 - Lifecycle monitoring – drift, performance and cost in production.
 - Responsible-AI review embedded in the delivery process.
 - ISO/IEC 42001 alignment formalised and internally reviewed.
 - Re-assessment – target overall ≥ 3.0 (Defined), no domain below 2.5.

NEXT STEP

Let's walk through this together.

This report is the start of a conversation. We'll talk through the findings, pressure-test the roadmap against your real constraints, and agree where to begin – no obligation.

[Book a readout →](#)

hello@millwater.consulting
millwater.consulting