

BEYOND THE INFRASTRUCTURE CHECKLIST

Building AI Readiness Across the Portfolio

A White Paper on Building Genuine AI Readiness Across the Enterprise Portfolio

Giovanni Leonardi

July 2023



Contents

1. Executive Summary	3
2. The Readiness Illusion	3
3. The Five Dimensions of AI Readiness	4
3.1 Dimension 1: Data Maturity	4
3.2 Dimension 2: Organisational Design	5
3.3 Dimension 3: Change and Adoption Capability	5
3.4 Dimension 4: Governance Readiness	6
3.5 Dimension 5: Technology Infrastructure	6
4. The Evidence: Why Infrastructure-Led Readiness Fails	7
4.1 Pattern 1: The Data Desert	7
4.2 Pattern 2: The Adoption Cliff	7
4.3 Pattern 3: The Governance Vacuum	7
5. A Readiness Framework for Portfolio Leaders	7
5.1 The Assessment Structure	8
5.2 The Readiness Sequence	8
6. Recommendation	8
7. Conclusion	9

Building AI Readiness Across the Portfolio

Portfolio · No. 93

Originally written July 2023 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“AI readiness is not a technology state — it is an organisational condition, and the organisations that assess it solely through the lens of infrastructure are the ones least likely to achieve it.”

1. Executive Summary

The rush to adopt artificial intelligence across enterprise portfolios has exposed a fundamental misdiagnosis: most organisations assess AI readiness as a technology question when it is, in fact, an organisational one. The pattern is consistent. An AI readiness assessment is commissioned. It examines cloud infrastructure, compute capacity, data lake architecture, and tooling. It concludes that the organisation is ready — or nearly ready — to deploy AI at scale. And then the programme fails, not because the infrastructure was inadequate, but because the data was unusable, the operating model was incompatible, the workforce was unprepared, and governance was absent.

This paper argues that AI readiness must be understood across five interdependent dimensions — data maturity, organisational design, change and adoption capability, governance readiness, and technology infrastructure — and that infrastructure is the least predictive of the five. It presents evidence from observed enterprise patterns, proposes a readiness framework that portfolio leaders can apply across their AI investments, and recommends a sequenced approach that addresses the hardest dimensions first.

2. The Readiness Illusion

The generative AI wave of 2022–2023 has accelerated an already dangerous tendency: the conflation of technical capability with organisational readiness. Large language models are more accessible than any previous AI technology. Cloud providers offer managed ML services that reduce the infrastructure barrier to near zero. The perception this creates — that AI is now easy to deploy — masks the reality that the organisational conditions for successful AI adoption have not changed at all.

The pattern I have observed across multiple sectors is remarkably consistent. An organisation conducts an AI readiness assessment. The assessment is led by the technology function. It evaluates the cloud platform, the data engineering pipelines, the MLOps toolchain, and the availability of data science talent. It produces a maturity score,

typically on a five-point scale, and recommends investments to close the gaps. The portfolio then proceeds on the assumption that readiness is a technical state that can be purchased.

Within twelve to eighteen months, the pattern reveals itself. AI proofs of concept succeed in controlled environments but fail to reach production. Models that perform well on curated datasets degrade on real operational data. Business units resist adoption because the AI outputs do not align with their workflows. Regulators ask questions the organisation cannot answer. The portfolio delivers activity — pilots, prototypes, demonstrations — but not value.

The readiness illusion is not that organisations overestimate their technical capability. It is that they fail to assess the four dimensions that actually determine whether AI will deliver value: data maturity, organisational design, change capability, and governance.

3. The Five Dimensions of AI Readiness

Genuine AI readiness is a condition of the whole organisation, not its technology function. It spans five dimensions, each of which must reach a minimum threshold before AI investments at portfolio scale can reliably deliver value.

3.1 Dimension 1: Data Maturity

Data is the most commonly cited prerequisite for AI and the most commonly neglected. The gap is not in data volume — most enterprises have vast quantities of data. It is in data quality, data governance, and data accessibility.

What readiness looks like:

- Data quality is measured, reported, and actively managed. There are defined data quality standards, automated quality checks, and clear ownership of data quality at source.
- Data governance is operational, not aspirational. Data owners are identified and accountable. Data lineage is documented. Access controls are enforced. Data catalogues exist and are maintained.
- Data is accessible to the teams that need it, in formats they can use, within timeframes that support iterative model development. The organisation has solved — or is actively solving — the data silo problem.
- Training data can be assessed for bias, representativeness, and fitness for purpose. The organisation can answer the question: does this data fairly represent the population the model will serve?

What readiness does not look like:

- A data lake exists but nobody trusts the data in it. Business users maintain their own spreadsheets because the official data sources are unreliable.

- Data governance is a policy document, not an operating discipline. Data owners are named but not active. Data quality is nobody's job.
- Data science teams spend 70–80 per cent of their time on data wrangling rather than model development — a figure that has persisted across the industry for a decade and shows no sign of improving.

In my experience, data maturity is the single largest determinant of AI programme success, and the dimension most often excluded from readiness assessments led by the technology function.

3.2 Dimension 2: Organisational Design

AI does not deploy into an organisational vacuum. It deploys into operating models, workflows, decision structures, and accountability frameworks. If the operating model is not designed to accommodate AI-assisted decision-making, the technology will be rejected by the organisation's immune system.

What readiness looks like:

- Roles and responsibilities have been redesigned to incorporate AI-assisted workflows. People know what the AI does, what they do, and where the boundary sits.
- Decision rights are clear: which decisions the AI informs, which it automates, and which it cannot touch.
- The operating model anticipates the change in work patterns that AI introduces — not just the elimination of tasks, but the creation of new ones (model monitoring, output validation, exception handling).

What readiness does not look like:

- AI is deployed as a bolt-on to existing processes with no redesign of workflows or roles. The AI produces outputs that nobody is accountable for acting on.
- The organisation treats AI adoption as a technology deployment rather than an operating model change.

3.3 Dimension 3: Change and Adoption Capability

Every AI deployment is a change management challenge. The technology changes how people work, what decisions they make, and what skills they need. Organisations without mature change capability will fail to adopt AI regardless of how good the technology is.

What readiness looks like:

- The organisation has a track record of managing technology-driven change — not just deploying systems, but changing how people work with them.
- Change management is funded, staffed, and treated as a programme workstream, not an afterthought.
- The workforce has been engaged early: expectations are managed, fears are addressed, and the value proposition is articulated in terms that matter to the people whose work will change.
- There is a realistic assessment of the capability gap between the current workforce and the workforce needed to operate in an AI-augmented environment, with a plan

to close it.

What readiness does not look like:

- Change management consists of a communications plan and a training schedule created three weeks before go-live.
- The organisation has a history of deploying technology and declaring victory when it goes live, rather than when people actually use it effectively.

3.4 Dimension 4: Governance Readiness

AI systems require governance structures that most organisations do not yet have. This is not about AI ethics policies — it is about the operational governance mechanisms that ensure AI systems are deployed responsibly, monitored effectively, and withdrawn when necessary.

What readiness looks like:

- There is a governance framework for AI that integrates with existing risk management and is not a standalone artefact.
- Decision-making authority is clear: who approves AI deployments, who monitors them in production, who can withdraw them.
- The organisation can explain its AI decisions to regulators, customers, and the public when asked.
- There is an incident response capability for AI failures — not a theoretical plan, but a tested process.

What readiness does not look like:

- AI governance is an ethics statement on the corporate website, unconnected to how AI systems are actually developed and deployed.
- Nobody can articulate who is accountable when an AI system produces a harmful or biased outcome.

3.5 Dimension 5: Technology Infrastructure

Technology infrastructure is the dimension most commonly assessed and the least predictive of success. This is not to say it is unimportant — adequate infrastructure is necessary. But it is the most commoditised, the most purchasable, and the least likely to be the binding constraint.

What readiness looks like:

- Cloud or on-premises compute is sufficient for model training and inference at the required scale.
- MLOps capabilities exist to support model versioning, deployment, monitoring, and rollback.
- Integration pathways exist between AI systems and the operational systems they must connect to.
- Security and access controls are appropriate for the data and models in use.

Most large enterprises are at or near readiness on this dimension. It is the easiest to assess, the easiest to remediate, and the least likely to be the reason an AI programme

fails.

4. The Evidence: Why Infrastructure-Led Readiness Fails

The case against infrastructure-led readiness assessment is not theoretical. The pattern of failure is observable across sectors.

4.1 Pattern 1: The Data Desert

Organisations that passed infrastructure readiness assessments with high scores have consistently failed when models encountered real operational data. Training data, curated for proof-of-concept demonstrations, bore little resemblance to the messy, incomplete, inconsistently labelled data that flows through production systems. The infrastructure was ready. The data was not.

4.2 Pattern 2: The Adoption Cliff

AI systems that performed well technically were rejected by the business users they were designed to serve. In financial services, fraud detection models that produced superior results to existing rules-based systems were overridden by analysts who did not trust them, did not understand them, and had not been involved in their design. In healthcare, clinical decision support tools sat unused because clinicians' workflows had not been redesigned to accommodate them. The infrastructure was ready. The organisation was not.

4.3 Pattern 3: The Governance Vacuum

AI systems were deployed without adequate governance, and the organisation discovered this only when something went wrong — a biased outcome, a regulatory inquiry, a press story. The remediation cost — both financial and reputational — vastly exceeded what governance investment would have cost upfront. The infrastructure was ready. The governance was not.

“AI readiness is not a technology state — it is an organisational condition, and the organisations that assess it solely through the lens of infrastructure are the ones least likely to achieve it.”

5. A Readiness Framework for Portfolio Leaders

The following framework provides a structured approach to assessing AI readiness across all five dimensions. It is designed for portfolio leaders, transformation directors, and CIOs

who need to make honest assessments of their organisation's readiness before committing to AI investments at scale.

5.1 The Assessment Structure

For each of the five dimensions, assess the organisation against four levels:

Level	Description
1 – Absent	No meaningful capability exists in this dimension
2 – Emerging	Some capability exists but it is fragmented, inconsistent, or dependent on individuals
3 – Established	Capability is defined, resourced, and operating consistently across the areas where AI is being deployed
4 – Optimised	Capability is mature, continuously improving, and actively enabling AI value creation

The minimum threshold for portfolio-scale AI investment is Level 2 across all five dimensions, with at least Level 3 in Data Maturity and Governance Readiness. Organisations below this threshold should invest in readiness before investing in AI.

5.2 The Readiness Sequence

The dimensions are not independent, and the order in which they are addressed matters. The recommended sequence is:

1. Data Maturity – first, because every other dimension depends on it. Without reliable, governed, accessible data, no AI investment will deliver sustained value.
2. Governance Readiness – second, because governance provides the framework within which all AI activity operates. Deploying AI without governance is accumulating risk that compounds over time.
3. Organisational Design – third, because the operating model determines whether AI outputs are used, not just produced.
4. Change and Adoption Capability – fourth, because change capability enables the workforce to transition to AI-augmented ways of working.
5. Technology Infrastructure – last, because it is the most commoditised and the most straightforward to remediate. It should not be neglected, but it should not lead.

This sequence is counterintuitive for many organisations, which naturally gravitate to infrastructure because it is tangible, purchasable, and within the technology function's direct control. But the sequence reflects the evidence: the dimensions that are hardest to build are the ones that matter most.

6. Recommendation

Portfolio leaders preparing their organisations for AI adoption should take three immediate actions:

1. Commission a five-dimension readiness assessment that covers data maturity, organisational design, change capability, governance, and technology infrastructure — not a technology-only assessment. The assessment should be conducted by a cross-functional team, not the technology function alone, and it should be brutally honest about the current state.
2. Invest in the foundations before the applications. Where the assessment reveals gaps in data maturity or governance, redirect investment from AI application development to foundational capability building. This is a difficult message to deliver in an environment where competitive pressure demands immediate AI results, but it is the only path to sustainable value. An AI application built on weak data and absent governance will fail — the only question is whether it fails visibly and quickly, or invisibly and expensively.
3. Reframe the portfolio narrative from AI deployment to AI readiness. The board and executive team must understand that readiness is the precondition for value, and that readiness cannot be purchased as a technology platform. It must be built as an organisational capability, across all five dimensions, and this takes time. The organisations that accept this reality and invest accordingly will be the ones that realise lasting value from AI. The organisations that skip the foundations will continue to produce impressive proofs of concept that never reach production, and impressive strategies that never reach reality.

7. Conclusion

The question are we ready for AI? is asked frequently across enterprise leadership teams. The answer, in most cases, is no — not because the technology is beyond reach, but because the organisational conditions for AI success have not been built. AI readiness is not a technology state. It is an organisational condition that spans data, governance, design, change capability, and infrastructure. Until portfolio leaders assess readiness across all five dimensions and invest in the hardest ones first, the gap between AI ambition and AI value will persist — and the growing catalogue of failed AI programmes will continue to expand.