



The Roadmap to Readiness

A CEO's guide to competing in an
AI-driven economy.

Table of contents

3	Introduction
4	Why AI is a CEO-Level issue
6	The Diagnostic Framework
8	Assessing AI Readiness
10	Assessing AI Impact
10	Introducing the Diagnostic Model
12	Determining How Fast to Move
15	From Readiness to Transformation
18	Conclusion

Introduction

Most leadership teams don't have a shortage of opinions about AI. They have a shortage of clarity about what to actually do with it.

The conversations are happening - in board meetings, in leadership offsites, in the informal exchanges between executives who are all quietly wondering whether their organisation is moving at the right pace. The pressure is real. Competitors are experimenting. Employees are already using tools, whether policies are in place or not. Investment cases are being made, pilots are being launched, and licences are being purchased.

And yet, for many organisations, the honest answer to “do we have a coherent AI strategy?” remains somewhere between “we’re working on it” and “not really.”

That gap, between the recognition that AI matters and the clarity about how to respond, is where a significant amount of time, money, and organisational energy is currently being lost. And it tends to produce one of two failure modes: the organisation that moves cautiously, building up readiness debt while competitors build capability; or the organisation that moves quickly, running up expensive pilots that generate enthusiasm but don't scale.

This whitepaper is written for leaders navigating that gap. It sets out a practical framework for assessing organisational readiness honestly, understanding competitive exposure in your specific context, determining the right pace of response, and sequencing the work that follows in a way that builds durable advantage rather than just activity.

The goal isn't to add to the noise. It's to help you cut through it.

Bimal Shah, Group CEO

Why AI is a CEO-Level Issue

AI is not simply the latest wave of digital transformation.

Previous technology shifts largely automated what people were already doing - faster, cheaper, at greater scale - whereas AI does something categorically different. It introduces a capability that can replicate, augment and in some cases replace cognitive work.

That distinction matters enormously for leaders, because it means the implications don't stop at the IT department, or even at operations. They reach into cost structures, workforce design, competitive positioning and margin, and they do so in ways that are still unfolding, which makes them harder to plan for and easier to underestimate.

To make this concrete: organisations are already seeing AI perform analysis, build financial models, write and review code, triage technical problems, and manage multi-step processes with minimal human intervention. The question is no longer whether this is coming. It's whether your organisation is positioned to understand and shape how it will happen and what impact it will have.

The economic consequences are worth naming directly. Cost structures will shift as some functions become increasingly AI-assisted or AI-led. Productivity gains are likely, but they don't automatically translate into margin improvement, particularly when AI-native competitors are using the same tools to commoditise what you've spent years building. Competitive advantage will increasingly belong to organisations with proprietary data, differentiated business models, and the operational capability to deploy AI with intent; not simply those with access to the best tools, which will soon be available to everyone.

None of this sits comfortably within a technology upgrade cycle. It is a strategic and operating model question - which means it belongs at the top of the organisation.

The CIO matters enormously in making AI work, but the decisions about where to invest, what to protect, how fast to move, and what kind of organisation you're building for the next decade? Those are CEO decisions. And organisations that treat them as anything less will find that out the harder way.

The Common Mistake: Ambition Without Diagnosis

The most common response to AI pressure looks something like this: a few pilots get launched, some enterprise licences get purchased, and employees are encouraged to “start using AI.” Copilots and generative assistants get rolled out. Basic guardrails are introduced to reduce the most obvious data risks. And then, broadly, everyone waits to see what happens.

This is understandable – and in many cases it’s a reasonable starting point. But it tends to produce a predictable set of problems. Experimentation becomes fragmented; teams grow uncertain about what’s actually permitted; pilots generate enthusiasm but stall before they scale; tool usage habits form that may not be privacy or data compliant; and the underlying questions about data quality, governance, and process redesign quietly get deferred until they become somebody else’s crisis.

Good Strategy Bad Strategy

We believe Richard Rumelt’s work “Good Strategy, Bad Strategy” is worth a mention here, as it describes exactly the discipline that AI demands of leadership teams right now.

Rumelt defines good strategy as “a cohesive response to an important challenge” – not a goal, not an aspiration, but a way through a specific difficulty.

Rumelt states that there are three key tenets to constructing good strategy:

- 1 A diagnosis** that defines or explains the nature of the challenge. A good diagnosis simplifies the often overwhelming complexity of reality by identifying certain aspects of the situation as critical.
- 2 A strategic response** for dealing with the challenge. This is an overall approach chosen to cope with or overcome the obstacles identified in the diagnosis.
- 3 A set of coherent actions** that are designed to carry out the strategic response. These are planned, coordinated steps enabling us to accomplish the strategic response together.

In constructing their strategic response, CEOs need to carry out a diagnosis in a way that will give them good insights into how they should start responding and how urgently they need to start this response.

What's Missing Isn't Activity; It's Diagnosis.

What's missing isn't activity; it's diagnosis. Most organisations investing in AI haven't yet asked the more fundamental question: what is AI actually doing to our competitive position, and are we genuinely prepared to respond? Declaring an ambition to "become AI-enabled" is not the same as having a strategy. And confusing the two is where a significant amount of AI investment quietly disappears.

The organisations that get this right don't necessarily move faster; rather, they move with more clarity about where AI creates real value for them specifically, what needs to be true before they can capture it, and what the cost of getting it wrong actually looks like.

That clarity starts with an honest diagnostic, not a tool rollout.

The Diagnostic Framework

Readiness & Impact

A coherent AI response requires clarity on two distinct but deeply connected dimensions – one internal, one external.

The first is readiness: how genuinely prepared is the organisation to adopt, govern, and scale AI effectively? The second is impact: how significantly is AI likely to reshape the competitive landscape the organisation is operating in? Together, these two dimensions determine not just whether to act and where to focus first but also how urgently we need to move and the costs of getting it wrong.

Assessing AI readiness

AI readiness is not defined by the number of tools deployed or the size of the licence spend. It is defined by an organisation's actual capacity to integrate AI responsibly and sustainably into how it operates, and that capacity is more human than it is technical.

At a leadership level, this means more than being broadly supportive of AI. It means executives being able to articulate a clear AI ambition that connects to the wider business strategy; understanding both the opportunity and the risk with enough depth to make real decisions; and creating the accountability structures that mean someone is genuinely responsible when things go wrong – not just when they go right.

But leadership is only part of the picture. Our work with behavioural scientists has consistently surfaced something that technical assessments tend to miss: the most significant constraint on AI readiness is usually cultural, not infrastructural. Whether employees feel confident using AI tools, whether they trust the technology, whether they believe it's safe to experiment and fail – these are not soft considerations. They are the difference between an organisation that learns quickly and one that accumulates expensive pilots that quietly go nowhere.

This shows up in specific, measurable ways. How capable do people feel in their use of AI; not just technically, but in terms of genuine self-efficacy? Do they see AI as an opportunity or a threat? What do social norms within the organisation reward – cautious avoidance or active experimentation? And critically: is there genuine psychological safety around failure, or does the culture make it easier to say nothing went wrong than to share what did?

Data and technology infrastructure matter too, of course. AI does not compensate for fragmented systems or poor-quality data; it amplifies weaknesses. Governance and ethics need to be explicit – clear policies on acceptable use, transparency, bias, privacy, and intellectual property are prerequisites for confident adoption, not afterthoughts. But the organisations that treat readiness as purely a technical question tend to be the ones most surprised when their AI investments underperform.

“AI readiness is not defined by the number of tools deployed or the size of the licence spend. It is defined by an organisation’s actual capacity to integrate AI responsibly and sustainably into how it operates, and that capacity is more human than it is technical.”

It’s also worth noting that the nature of the readiness question will shift as AI matures. As AI becomes as embedded in working life as email, the question of whether people are using it becomes less interesting than whether they’re using it well – in their specific role, within a well-governed organisation, in ways that actually connect to outcomes. A readiness framework built only on today’s questions will need to evolve; one built on the underlying dimensions of leadership, culture, psychological safety, and organisational systems is considerably more durable.

Ultimately, readiness is a multi-dimensional leadership capability, not a technical checklist.

Assessing AI Impact

If readiness reflects what's happening inside the organisation, impact reflects exposure to what's happening outside it.

And the honest answer is that exposure varies enormously - which is precisely why generic urgency ("every organisation needs to act now") is less useful than a clear-eyed assessment of your specific situation.

Some industries are already experiencing significant AI-driven disruption, as new entrants deploy AI-native business models that structurally undercut established players on cost, speed, or personalisation. Others are seeing slower shifts - but still face meaningful pressure along specific parts of their value chain, even if the headline competitive picture looks stable.

The factors that determine exposure include the degree of standardisation in core processes, the intensity of cognitive work across the organisation, the pace of competitor adoption, and the evolving expectations of customers. In some sectors, customers are already expecting AI-enabled speed and personalisation and adjusting their behaviour accordingly. In others, trust and empathy remain central differentiators - and over-automation doesn't just fail to add value, it actively destroys it. Efficiency gains that erode customer trust are not strategic wins; they're margin improvements that borrow against reputation.

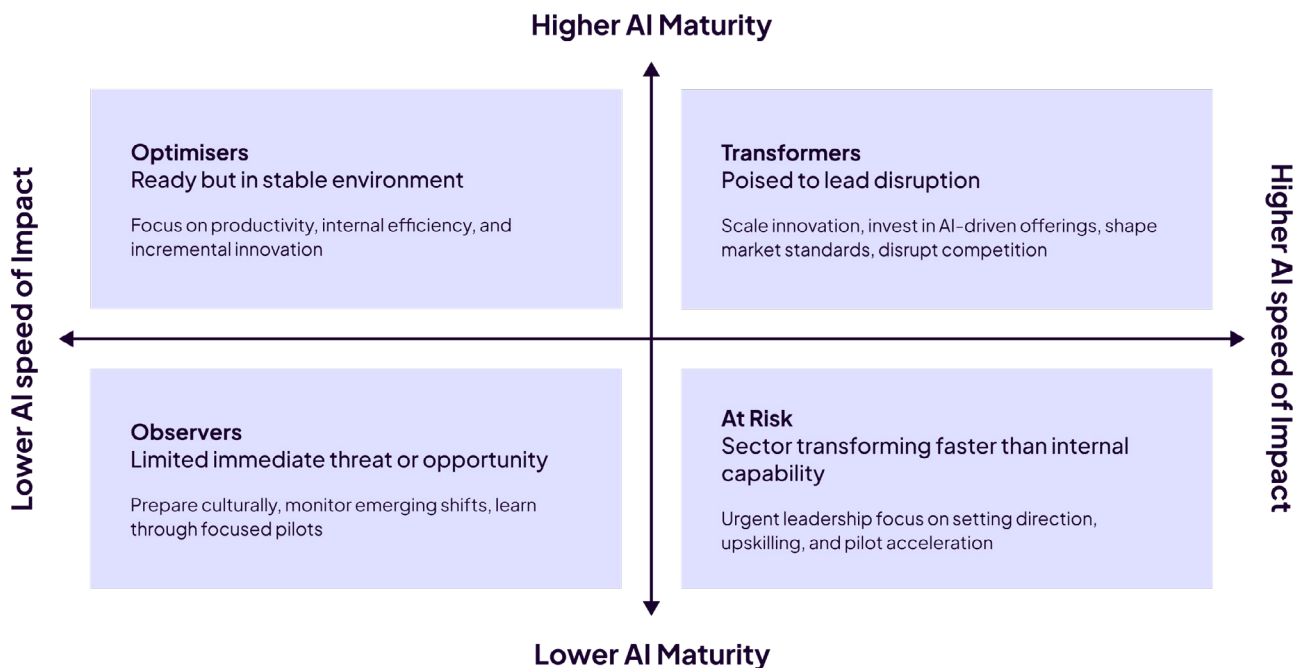
Introducing the Diagnostic Model

Readiness and impact together define the appropriate pace and shape of response - and the interaction between them is what matters most.

An organisation with low readiness but low competitive exposure has a window to build foundations thoughtfully. An organisation with low readiness and high competitive exposure has a more urgent problem. And the worst position of all is high exposure, low readiness, and a leadership team that doesn't have clarity on either.

Working with a team of behavioural scientists, we've developed a proprietary diagnostic that assesses both dimensions - combining qualitative and quantitative measures and designed to generate clear output without significant disruption to the organisation.

The diagnostic model places organisations across two axes, producing a view of where they sit today and what that means for their immediate priorities.



The diagnostic is not a score.

It's a starting point for a strategic conversation - one that maps clearly onto specific actions, because a tool that surfaces a problem and then leaves the organisation to figure out what to do about it is, frankly, not very helpful.

Determining How Fast to Move

The diagnostic places organisations into one of four positions – and each position implies a different kind of response.

1 At Risk

Sector transforming faster than internal capability

For organisations that are At Risk – highly exposed to AI disruption but poorly prepared to respond – the situation is straightforward, if uncomfortable. The window for building foundations thoughtfully is already narrowing. The priority is not transformation; it's getting the basics right quickly enough to avoid being structurally disadvantaged before the opportunity to catch up disappears.

2 Observers

Limited immediate threat or opportunity

For Observers – organisations with low readiness but operating in environments where competitive pressure is building steadily rather than suddenly – the opportunity is in disciplined, selective scaling. The foundations need to be established in a timely manner; the question is then where AI will create the most durable advantage, and how to capture it before competitors get there first.

3 Optimisers

Ready but in stable environment

For Optimisers – those with genuine readiness but lower immediate exposure – the risk is a different one. It's not acting too soon; it's failing to build and test capabilities during the window when the cost of experimentation is still low. The organisations that will be best positioned when disruption accelerates in their sector are the ones that didn't wait until it was already visible.

4 Transformers

Poised to lead disruption

For Transformers – strong readiness, high impact environments – the strategic question shifts. The foundations are in place, and the competitive stakes are clear; the focus now is on embedding AI into the core operating model in ways that are difficult to replicate, and on staying ahead of the economic and workforce implications as AI capability continues to develop.

In almost every case, the most dangerous position is the one that doesn't feel dangerous – the organisation that has quietly decided to wait, without making a conscious decision.

There's a meaningful difference between a deliberate choice to move at a measured pace, grounded in a clear understanding of readiness and exposure, and an absence of decision that's accumulated into a posture. The former is something you can defend in a board meeting, whereas the latter tends to become apparent only in retrospect, and usually at some cost.

What good strategy looks like here is not complicated, even if the execution is: an honest diagnosis of where you actually sit, a coherent response that reflects that reality, and the organisational alignment to act on it.

The organisations that get this right won't necessarily be the fastest movers. But they will be the most intentional ones, and in a domain where the stakes are rising, that distinction matters more than it might appear.

“The roadmap to readiness isn’t about slowing down.

It’s about making sure that when you accelerate – and you will probably need to – you’re doing so with a clear view of the road ahead, rather than hoping the ground is solid enough to hold.”

- The Roadmap to Readiness

From Readiness to Transformation

Understanding where you sit on the diagnostic is the starting point, not the destination.

The more important question is what you do with that clarity - and how you sequence the work that follows. Because AI transformation doesn't happen in a single initiative or a big-bang deployment; it unfolds in stages, and the organisations that navigate it well are the ones that resist the temptation to skip the unglamorous early work in favour of the more visible stuff.

That early work starts with foundations - not in the abstract sense of "getting ready to get ready," but in the practical sense of establishing what needs to be true before AI investment can generate reliable returns. Strategic intent needs to be clear enough that people across the organisation understand what AI is for, and what it isn't. And the pilots that get run at this stage should be chosen for what they teach the organisation, not for how well they're likely to look in a press release.

The policy piece comes first.

Before any AI implementation gets underway, the organisation needs clear guardrails covering what AI can be used for, by whom, with what data, and under what conditions.

Teams that don't know what's permitted tend to either avoid AI entirely or use it in ways that create risk - neither of which is the outcome anyone is looking for.

Then comes the data.

The question most organisations should be asking isn't "do we have enough data?" - it's "is our data actually fit for what we're asking AI to do with it?"

Those are very different questions, and the gap between them is where a lot of AI investment quietly goes to work without producing much in return.

A useful assessment looks at four things:

- 1 Whether the data is structured and documented correctly for the intended use
- 2 Whether it's actually the right data for the outcome
- 3 Whether the organisation is legally and contractually permitted to use it in the way intended
- 4 Whether it's complete and representative enough to produce outputs that can be trusted, without gaps that introduce bias or blind spots.

Outputs that can't be trusted don't get used, which tends to be a difficult conversation to have after the investment has already been made.

Data readiness is a leadership responsibility, not a technical one – and the compliance and regulatory environment sits within the same frame.

GDPR, intellectual property, sector-specific requirements: these need to be part of the conversation from the start, not discovered as a constraint halfway through a deployment.

We've explored both in considerably more depth elsewhere; if either is on your radar (and it should be), the reading is worth your time.

With strategy set, policy established, and data assessed, the organisation is genuinely ready to move into the build, test, and iterate process – developing a proof of concept that reflects real conditions rather than optimistic assumptions, and learning from it in a way that compounds over time.

The sequencing matters more than the ambition. Attempting to shortcut this process doesn't accelerate progress; it creates instability, erodes trust, and produces the kind of wasted investment that makes the next AI initiative harder to fund and harder to believe in.

Readiness before acceleration.

AI will reshape how organisations create value and compete; that much is not seriously in dispute.

The more interesting question, and the one that actually separates the organisations that will come out of this period well from those that won't, is whether leaders have an accurate picture of where they stand; not where they'd like to be, or where they assume they are, but where they actually are. Most don't. And that gap, between assumed readiness and actual readiness, is where AI investment has a habit of disappearing without much to show for it.

The leaders who navigate this well won't necessarily be the ones who move fastest or spend most. They'll be the ones who took the time to understand their competitive exposure honestly, assessed their organisational readiness without flattering themselves, and made deliberate decisions about pace and focus as a result. That's not caution for its own sake - it's what good strategy has always looked like, applied to a domain where the stakes are rising and the cost of getting it wrong is increasingly hard to ignore.

AI readiness is not a project with a completion date. It's an ongoing leadership capability - one that touches strategy, people, data, governance, operating model, and economics simultaneously, and that needs to be actively maintained as the technology, the competition, and the regulatory environment continue to evolve.

The roadmap to readiness isn't about slowing down. It's about making sure that when you accelerate - and you will probably need to - you're doing so with a clear view of the road ahead, rather than hoping the ground is solid enough to hold.

About Purple Shirt

Purple Shirt is a customer experience, AI, and technology consultancy with deep roots in human centred design. We help organisations make better choices, and build the products, services, and capabilities that turn those choices into outcomes.

We were recently acquired by Elemental Concept, a global technology consultancy bringing together leaders in AI, customer experience, change, and engineering. The combined group gives our clients access to senior expertise across the full arc of a digital or AI transformation, from setting direction through to working software in production.

Our services span four connected areas:

Strategy and innovation

Setting direction, identifying where value is created, and sequencing the work that follows.

Human centred experience design

Designing services and products around the people who use them, grounded in research.

AI, data and analytics

Building the data foundations and AI capabilities that turn information into advantage.

Technology and engineering

Designing, building, and operating the systems that make the strategy real.

Find out more
purpleshirt.co.nz

Get in touch
steve@purpleshirt.co.nz

