

ADVISORY BRIEF

Building AI Capability Takes Time. That's Why Starting Now Is an Advantage.

AI adoption is happening faster than any prior technology wave. The potential is real. Organisations that start building capability now have the luxury of learning at their own pace. Here's a staged approach.

Building AI Capability Takes Time. That's Why Starting Now Is an Advantage.

March 2026 · For CEO, CIO, Board, COO

EXECUTIVE SUMMARY

AI adoption is happening faster than any prior technology wave. The potential is real. Organisations that start building capability now have the luxury of learning at their own pace. Here's a staged approach.

CONTENTS

1. The Adoption Curve Is Familiar. The Speed Is Not.	4
2. The Customer Expectation Shift Worth Preparing For	4
3. What Early Adopters Are Discovering	5
4. Common Reasons for Delay, and Why Experimentation Addresses Most of Them	6
5. The Case for a Staged, Measured Approach	6
6. The Value of Learning Time	7

Building AI Capability Takes Time. That's Why Starting Now Is an Advantage.

Advisory Brief · March 2026 · For CEO, CIO, Board, COO

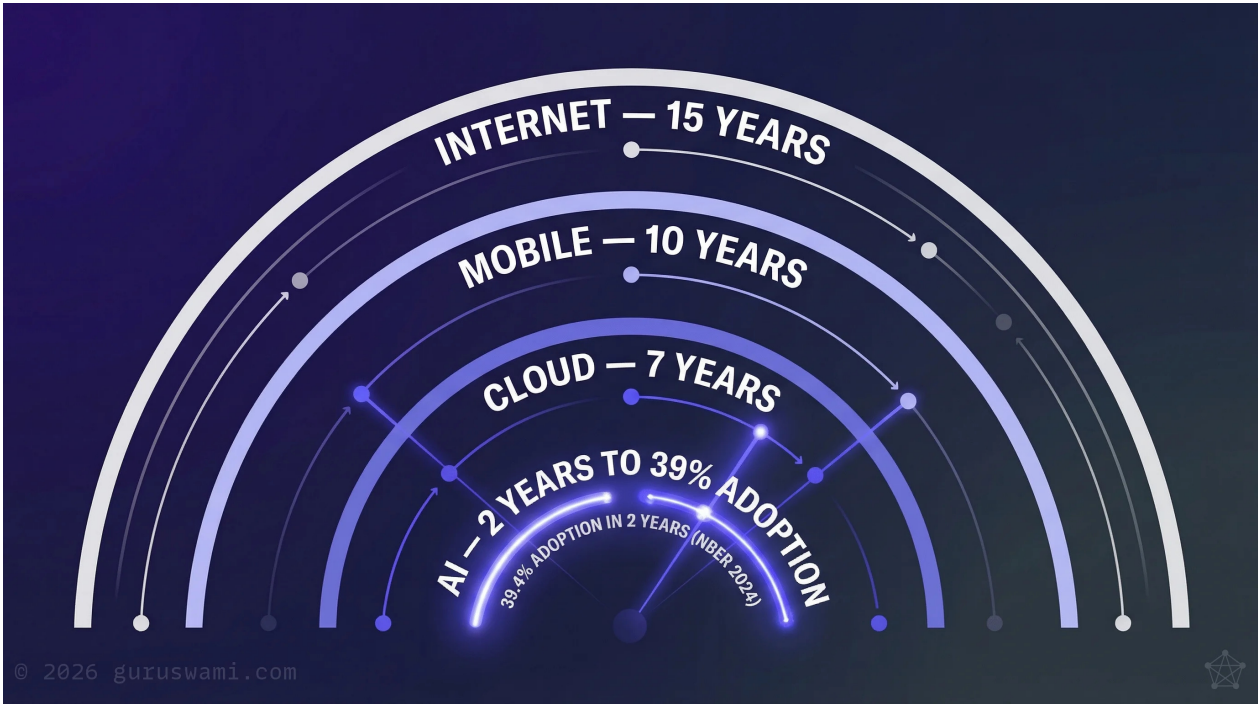


Exhibit 1 The AI Adoption Clock, technology adoption timelines compressing from 15 years to 3-5 years

Source: Guruswami Advisory

39.4%

adoption in just two years. ^[1] The internet and PCs were at 20% at the same point. Generative AI is being adopted at roughly double the speed of any prior technology wave. **Organisations building capability now have time to learn at their own pace.**

Federal Reserve Bank of St. Louis / NBER, 2024

KEY ARGUMENTS IN THIS BRIEF

1. **AI adoption is faster than any prior technology wave.** 39.4% adoption in two years, compared to 20% for PCs and the internet at the same stage.
2. **Building capability takes time.** The real advantage is depth of understanding, not speed of purchase. Institutional knowledge compounds through months of hands-on work.
3. **Common reasons for delay do not survive scrutiny.** The technology is deployable, the regulatory trajectory is clear, and the business case builds through experimentation, not observation.
4. **A staged approach reduces risk.** Start with internal capability on local infrastructure, then expand through controlled deployment, customer-facing AI, and emerging capabilities.
5. **Starting now is a luxury, not a pressure.** Organisations that begin building capability while the cost of mistakes is low will have the judgment others wish they had when the pace accelerates.

The Adoption Curve Is Familiar. The Speed Is Not.

Every transformative technology follows a recognisable adoption curve: early experimentation, accelerating expectations, broader uptake, and eventually a new baseline. We saw it with the internet, with mobile, with cloud computing.

AI is following the same pattern. The difference is speed. The gap between "interesting experiment" and "standard expectation" is shorter than it was for any prior wave.

That speed is worth understanding, because it tells you something important about the value of starting early. Building real AI capability, the kind where your people understand the technology well enough to use it safely, takes time. Months of hands-on work, iteration, shared vocabulary. Starting now means doing that learning while you have room to think.

The Customer Expectation Shift Worth Preparing For

Right now, most customers don't expect a fully interactive, 24/7 speech-to-speech helpdesk that can respond to questions, process orders, and resolve problems in real time. They're still accustomed to hold queues, business-hours-only support, and FAQ pages.

By 2027, they may be disappointed not to have it.

The technology already exists. Real-time voice agents capable of natural conversation, with access to customer records, order systems, and knowledge bases, are deployable today. Not as a research prototype. As production infrastructure. Organisations moving on this now are building the operational muscle to do it well. They are training their teams, refining their data pipelines, understanding the failure modes, and learning how to govern AI that speaks directly to customers.

Organisations building this capability now are doing it with time to think, experiment, and learn from mistakes. That institutional knowledge is the real advantage. Deploying AI-powered customer interactions without adequate testing, governance, or team capability is where the real risk sits.

What Early Adopters Are Discovering




Exhibit 2 What early adopters are discovering: operational efficiency, faster iteration, better service design, and readiness for emerging capabilities

Source: Guruswami Advisory

Organisations that started building AI capability early are reporting tangible results. Not because they rushed. Because they had time to learn.

Operational efficiency. 54% of infrastructure leaders are already adopting AI specifically to reduce costs^[2] (Gartner, 2025). These are not speculative pilots. They are measured deployments where teams understood the technology well enough to target genuine cost drivers.

Faster iteration. AI-augmented product development, market analysis, and decision-making compress timelines. Organisations that have built this into their workflows report that the time savings come not from the technology alone, but from their teams knowing how to use it well.

Better service design. When teams understand what AI can and cannot do, they design better customer experiences. They know where human oversight is needed. They know which failure modes to test for. This knowledge only comes from practice.

Readiness for what comes next. AI is generating new categories of capability: applications and business models that don't exist yet but will emerge. Organisations with internal AI competence will be positioned to evaluate these as they appear. On their own terms and timeline.

Common Reasons for Delay, and Why Experimentation Addresses Most of Them

Most organisations delay for reasons that feel rational. They wait for the technology to settle, for regulatory clarity, for a proven business case. Each of these is understandable.

But the technology is already deployable. The [regulatory trajectory](#)^[3] is clear: the [DTA](#)^[4], [APRA](#)^[5], and [ASIC](#)^[6] are all moving toward mandatory requirements, not away from them. And the business case builds through experimentation, not observation. The governance concerns are real, and they are solvable. We cover [data exposure](#)^[7] and [agentic attack surfaces](#)^[8] in separate briefs. This one is about how to move anyway.

The good news is that none of these concerns require waiting. Low-risk internal pilots on local infrastructure let teams learn without data sovereignty exposure, without metered costs, and without regulatory risk. The clarity most organisations are waiting for is the kind that comes from doing the work, not from watching others do it.

The Case for a Staged, Measured Approach

The answer is not to rush. Reactive adoption is precisely how organisations introduce the risks they were trying to avoid. Data breaches from hasty integration. Hallucinating AI systems facing customers. Regulatory non-compliance from ungoverned deployments. Teams who don't understand the technology they're operating.

The better path is to start now, with a staged approach that builds capability ahead of need.



Exhibit 3 Staged AI Adoption Roadmap: 4 phases from internal capability through controlled deployment, customer-facing AI, and emerging capabilities

Source: Guruswami Advisory

Phase 1: Internal capability (now). Build the team's understanding of the technology through hands-on experimentation. Local inference environments allow unlimited learning without metered costs or data sovereignty concerns. Start with concrete tasks: summarising risk reports, drafting internal policies, triaging support tickets. Identify where AI has the most practical impact on your specific operations.

Phase 2: Controlled deployment (2026). Move proven use cases into production with proper governance, monitoring, and human oversight. Start with internal-facing applications before customer-facing ones. Establish measurement frameworks that demonstrate ROI to the board.

Phase 3: Customer-facing AI (2026 to 2027). Deploy AI-powered interactions with the confidence that comes from institutional knowledge, not vendor promises. Teams understand the failure modes because they've encountered them in controlled environments. Governance frameworks exist because they were built during earlier phases.

Phase 4: Emerging capabilities (ongoing). With a mature internal capability, an organisation is positioned to evaluate and adopt new AI capabilities as they emerge. On its own timeline.

The Value of Learning Time

The real advantage of starting now is not speed. It is depth of understanding.

AI capability is built through practice. Teams that have spent months working with models, testing failure modes, refining data pipelines, and building governance frameworks have something that cannot be purchased or compressed: institutional knowledge. They know what works in their specific context. They know what fails. They know how to explain it to their board.

Every organisation will get there. Nobody is going to miss out. But starting now means your people learn while they have time to think, experiment, and iterate without pressure. That learning compounds. Each quarter of structured experimentation builds on the last.

The potential is real. A staged, measured approach is how you realise it safely.


KEY TAKEAWAYS

- AI adoption is happening faster than any prior technology wave. Building real capability takes time. Starting now means learning at your own pace.
- Organisations already building AI capability report real gains in efficiency, service quality, and decision speed. The gains come from understanding, not just tooling.
- Low-risk internal pilots on local infrastructure address most common reasons for delay: no data sovereignty exposure, no metered costs, no regulatory risk.
- A staged approach (internal capability, controlled deployment, customer-facing AI) builds the institutional knowledge to deploy confidently.
- Every organisation will incorporate AI. Starting earlier means more room to experiment, more room to make mistakes safely, and deeper understanding when it counts.

Guruswami Advisory helps Australian organisations build AI capability at their own pace. Staged, measured, grounded in practice. Every recommendation is tested on our own distributed inference lab before it reaches yours.

References

1. <https://www.stlouisfed.org/on-the-economy/2024/sep/rapid-adoption-generative-ai>
2. <https://www.gartner.com/en/newsroom/press-releases/2025-10-29-gartner-survey-54-percent-of-infrastructure-and-operations-leaders-are-adopting-artificial-intelligence-to-cut-costs>
3. <https://guruswami.com/insights/strategic-roadmap-2026/#pillar-iv-the-regulatory-calendar>
4. <https://www.digital.gov.au/ai/ai-in-government-policy>
5. <https://www.apra.gov.au/operational-risk-management>
6. <https://www.asic.gov.au/regulatory-resources/find-a-document/reports/rep-798-beware-the-gap-governance-arrangements-in-the-face-of-ai-innovation/>
7. <https://guruswami.com/insights/shadow-ai-data-leaving-your-network/>
8. <https://guruswami.com/insights/agent-ai-next-attack-surface/>



About Guruswami Advisory

Independent AI security and strategy advisory for Australian boards, leadership teams, and regulated organisations. No vendor ties. No platform allegiance. Every recommendation tested on our own infrastructure.

Paul Nevin, Principal Advisor. 28 years in cybersecurity and cyber-intelligence. Six years of applied AI research. Every engagement led personally.

Contact

info@guruswami.com

guruswami.com

[linkedin.com/in/paul-nevin](https://www.linkedin.com/in/paul-nevin)

Guruswami™ Pty Ltd | ABN 11 695 354 020 | Canberra, ACT, Australia

This document is provided for informational purposes. It does not constitute legal, financial, or insurance advice. Where findings have regulatory implications, engage qualified legal counsel.