




ADAPT

From AI Ambition to Execution Reality

Why AI outcomes now depend on data maturity,
operating discipline, and trusted execution

Australian organisations are no longer debating whether AI matters. The question now is whether the enterprise conditions required to scale it are actually in place.



Executive summary

Australian organisations have moved beyond AI curiosity. The strategic case for AI is now broadly accepted across executive teams. What is less settled is whether the organisational conditions required to scale AI are actually in place. ADAPT's latest research across CIOs, CDAOs, CFOs, infrastructure leaders, digital leaders, and government leaders points to the same conclusion: ambition has accelerated, but execution maturity remains uneven.

ADAPT's latest research indicates that the market has entered a different phase. Across nearly 600 CIOs surveyed over the past 12 months, AI strategy rose from the fifth-ranked priority to the top priority in 2026¹, overtaking broad-scale technology modernisation and cost optimisation. The issue is no longer whether AI matters, but what is preventing organisations from scaling it.

The evidence suggests that scale is being constrained less by intent than by enterprise conditions.




Data maturity
remains uneven



Integration and interoperability
are still lagging in many environments.



Cloud and platform investments
are progressing, but operating models, workflow redesign, and governance maturity are not keeping pace.



Organisations are moving forward,
but often in bounded, selective, and tightly supervised ways.

This marks a clear shift from last year's readiness-led conversation. The market has moved from asking "are we ready for AI?" to asking "what is stopping us from scaling it?" The answer increasingly lies in the interaction between data foundations, modernisation maturity, governance discipline, and delivery capacity. That is why this report focuses less on AI promise in the abstract, and more on the practical conditions that determine whether AI outcomes can be operationalised at scale.

The next phase of AI outcomes will be shaped less by access to tools and more by execution conditions. Scale is increasingly being determined by four deeper realities: the maturity of the data environment, the unevenness of modernisation progress, the growing influence of governance and assurance, and the increasing reliance on integrated platforms and trusted partners to offset capability and operating gaps.

1. ADAPT CIO Edge Surveys in Feb and Aug 2025 and Feb 2026, n=597

The operating shift

The market has changed. A readiness-led lens is no longer sufficient because the core challenge is no longer whether organisations are starting. While many already have, the harder question is whether they can scale AI while managing legacy complexity, fragmented data conditions, uneven modernisation, rising governance expectations, and limited internal capacity.

That shift matters because it changes how progress should be interpreted. In earlier stages of the market, AI ambition itself was a useful signal. Now it is insufficient. ADAPT's CIO work shows AI strategy has become the top priority², but the same research also shows that about 28% of mission-critical applications still depend on legacy platforms³. This suggests that many organisations are trying to scale new capabilities while still carrying a heavy burden of historical complexity.

The inflection point: AI ambition is surging, but execution is lagging

Top 10 goals 2H 2025

- 1 Tech modernisation and simplification
- 2 Improving operational effectiveness
- 3 Grow the business
- 4 Optimising costs
- 5 Developing the AI strategy and roadmap
- 6 Building a secure and trusted organisation
- 7 Acquiring and retaining customers
- 8 Creating a data-driven organisation
- 9 Pursue broad-based digital transformation
- 10 Digitisation and automation of key processes


Top 10 ranked goals for 2026

- 1 Developing the AI strategy and roadmap
- 2 Grow the business
- 3 Improving operational effectiveness
- 4 Tech modernisation and simplification
- 5 Optimising costs
- 6 Acquiring and retaining customers
- 7 Developing new products and services
- 8 Creating a data-driven organisation
- 9 Driving financial performance and value creation
- 10 Leading organisational change and transformation

2. ADAPT CIO Edge Surveys in Feb and Aug 2025 and Feb 2026. N = 597 Australian CIOs

3. ADAPT CIO Edge Survey in Feb 2026. N = 194 Australian CIOs

The same operating tension appears from the finance lens. ADAPT’s CFO research shows that, on average, only 61% of organisations’ current technology capabilities are adopted and in use, implying that 39% of technology spend is not converting into practical value⁴. That is an important backdrop for any AI narrative. It means the market is not approaching AI as a clean-sheet opportunity. It is approaching AI through an environment already shaped by underused capability, scrutiny on value, and tighter expectations around business justification.




“AI ambition is real. Investment is real. Modernisation is real. But scale is now being filtered through a more demanding execution lens.”

From AI intent to execution friction

The first major shift in the market is that AI ambition is no longer the binding constraint. Australian organisations have moved past curiosity, experimentation for its own sake, and high-level readiness discussions. The more useful question now is whether the enterprise can absorb AI into its operating environment without being slowed by data fragmentation, immature delivery practices, inconsistent governance, and unclear ownership.

ADAPT's CIO research shows a market that has moved beyond intent and into the harder terrain of execution. Among CIOs, 34% remain at initial awareness and 36% at active exploration of agentic AI. Only 22% report targeted implementation in selected teams, 8% report operational scaling, and none report value optimisation at scale⁵. Progress is visible, but the market remains early.


How mature is your organisation's overall use of agentic AI? CIO Aug 2025 vs Feb 2026



5. ADAPT CIO Edge Survey in Aug 2025 & Feb 2026. N = 361 Australian CIOs

The ROI picture is similarly constrained. When asked if you have seen measurable ROI from GenAI (copilots, model consumption, embeddings) yet, 76% of CDAOs reported either no measurable ROI to date or only early signs with limited or unclear returns. Only 19% said ROI was broadly meeting expectations, while 5% said it was exceeding expectations in some areas. Large enterprises are further ahead, but broad-based returns remain limited across every size segment. The dominant pattern across every segment is still early or uneven value realisation, showing that most organisations are still moving from experimentation to repeatable business impact.

Measurable ROI from GenAI (copilots, model consumption, embeddings)




Among CIOs, 84% report either no measurable ROI from agentic AI to date or only early signs with limited or unclear returns. Among CDAOs, the equivalent figure is 78%⁶. The question facing leaders is no longer whether AI belongs on the agenda. It is why momentum slows when ambition is already high.

6. ADAPT CIO, and Data and AI Edge Surveys in Feb and Apr 2026. N = 358 Australian CIOs and CDAOs

The structural ceiling

AI outcomes are increasingly being decided by structural conditions beneath the tooling layer. Data maturity and modernisation maturity are emerging as the primary determinants of what can scale, how quickly it can scale, and how safely it can be sustained.

Top CIO constraints to scale Agentic AI (% of respondents)



Respondent profile includes:


Large enterprise (>10,000 emp.)	19%
Enterprise (2501 to 10,000 emp.)	33%
Mid-market (501 to 2500 emp.)	35%
Small / medium corporate (1-500 emp.)	12%

Data foundations now decide how far AI can go

ADAPT's latest Data and AI research shows that the market is no longer merely asking which AI tools to buy. It is asking whether the underlying data and operating environment can sustain them. The maturity curve is still concentrated in the middle: 46% of CDAOs place their organisations in the emerging category and 41% in scaling, while only 6% report optimised maturity and 2% transformational maturity.

For mid-market and smaller organisations, this matters even more, as maturity remains concentrated in the emerging and scaling stages rather than in more advanced states, suggesting the next barrier to AI scale is less tool access than the strength of the underlying data environment.⁷

Where would you place your organisation on a Data & AI maturity curve?




7. ADAPT Data and AI Edge Survey in Apr 2026. N = 162 Australian CDAOs

The readiness picture becomes even tighter when the question shifts from broad maturity to data preparedness for AI at scale. Only 8% of organisations rate their data environment as optimised, and none rate it fully ready. Nearly half, 48%, still describe themselves as emerging, with another 35% developing. Notably, the share rating themselves optimised has fallen from 23% in the prior year to 8% in the latest cut⁸, suggesting that as expectations rise, more organisations are recognising how much foundational work remains.

How prepared is your organisation's data to enable AI at scale? Time Series

■ 2025 ■ 2026




8. ADAPT Data and AI Edge Survey in May 2025 and Apr 2026. N = 280 Australian CDAOs

Architecture and operating discipline are also limiting scale. When asked how well their current data and AI architecture supports scalability, modularity, and the ability to adopt new GenAI models without major re-engineering, only 17% of organisations described their architecture as highly capable. Another 41% see it as only partially capable, 31% as offering limited support, and just 3% as fully future-ready.⁹

For mid-market and smaller organisations, the picture is more constrained: only a small minority say their architecture is highly capable, while most continue to operate in environments that are either limited in support or only partially capable. Large enterprises are further ahead, with 44% reporting highly capable architectures. This suggests that, for much of the smaller-organisation market, architecture remains a practical constraint on how quickly AI can scale across the enterprise.

Data & AI architecture maturity to enable genai without major re-engineering

- Not at all (legacy/siloed architecture; major rebuild needed)
- Limited support (some modern components but scaling/integration difficult)
- Partially capable (supports moderate AI workloads; struggles with rapid model evolution)
- Highly capable (scalable, modular, good AI/ML/GenAI integration)
- Fully future-ready (cloud-native, composable, optimised for continuous model evolution)



9. ADAPT Data and AI Edge Survey in Apr 2026. N = 162 Australian CDAOs

On operational practices for deploying, monitoring, and maintaining AI and ML systems, only 12% rate themselves advanced or fully mature. The constraint has moved below the model layer. AI ambition is advancing faster than the enterprise’s ability to trust, connect, govern, and operationalise the data environment beneath it.

Modernisation progress is real, but unevenly translating into outcomes

The same pattern appears in platform and infrastructure decisions. Significant legacy dependence still sits inside many enterprises, and that load competes directly with capacity for new AI-enabled value. CIOs report that 27.8% of mission-critical applications still depend on legacy platforms. The issue is not lack of activity. Too much enterprise energy is still being absorbed by carrying and modernising the core.

The infrastructure view adds a second layer. ADAPT’s Cloud and Infrastructure research shows a market dealing with hybrid complexity, selective repatriation, and rising infrastructure demand rather than a simple march to cloud-native maturity. In 2025, 25% of infrastructure leaders say they are planning to repatriate public-cloud workloads, up from 10% in 2020. Among those repatriating, the expected average share brought back over the next 12 months is 21% of public-cloud workloads¹⁰.

Modernisation progress itself remains uneven. Infrastructure leaders estimate average progress at 44% for infrastructure modernisation and consolidation, 41% for API connectivity, 37% for workload interoperability across multi-cloud environments, 36% for workflow automation, and just 30% for application consolidation. Automation maturity tells a similar story: 22% still report limited automation and primarily manual processes, while only 20% are at comprehensive DevOps automation or full-scale GitOps. Modernisation is underway, but benefits diminish where integration and operating models fail to keep pace.

10. ADAPT Cloud and Infrastructure Edge Survey in July 2025. N = 96 Australian Infrastructure leaders


Security, governance, and assurance are the new gatekeepers

Security, governance, and assurance are no longer peripheral constraints on AI. They are becoming the boundary conditions of what can scale. ADAPT's latest research tracks who leads AI governance, who should lead it, how mature governance is, how confident organisations are in managing AI-related risks and regulation, and how much of the AI environment is formally covered by governance or risk frameworks. That evidence points to a market entering a more demanding phase.

The current picture is mixed. 61% of CIOs say AI governance currently sits with CIO or technology leadership¹¹. Yet governance maturity remains shallow: only 15% of respondents sit in the top two maturity categories. Coverage is also uneven. Just 29% say more than three quarters of AI pilots, tools, or models are known and governed, while 7% say AI use remains largely untracked and another 19% say fewer than a quarter of AI tools are known and governed.

Readiness at the workforce level is equally important. Only 15% of CIOs report mandatory AI awareness training for all staff, while 9% say they still have no formal training and no policies in place. The rest are clustering around lighter forms of preparedness: 30% have AI usage guidelines and policies, 27% are encouraging AI-assisted workflows with optional upskilling, and 19% have established AI champions or a centre of excellence¹². Governance is no longer only a control question. It is also an organisational preparedness question.

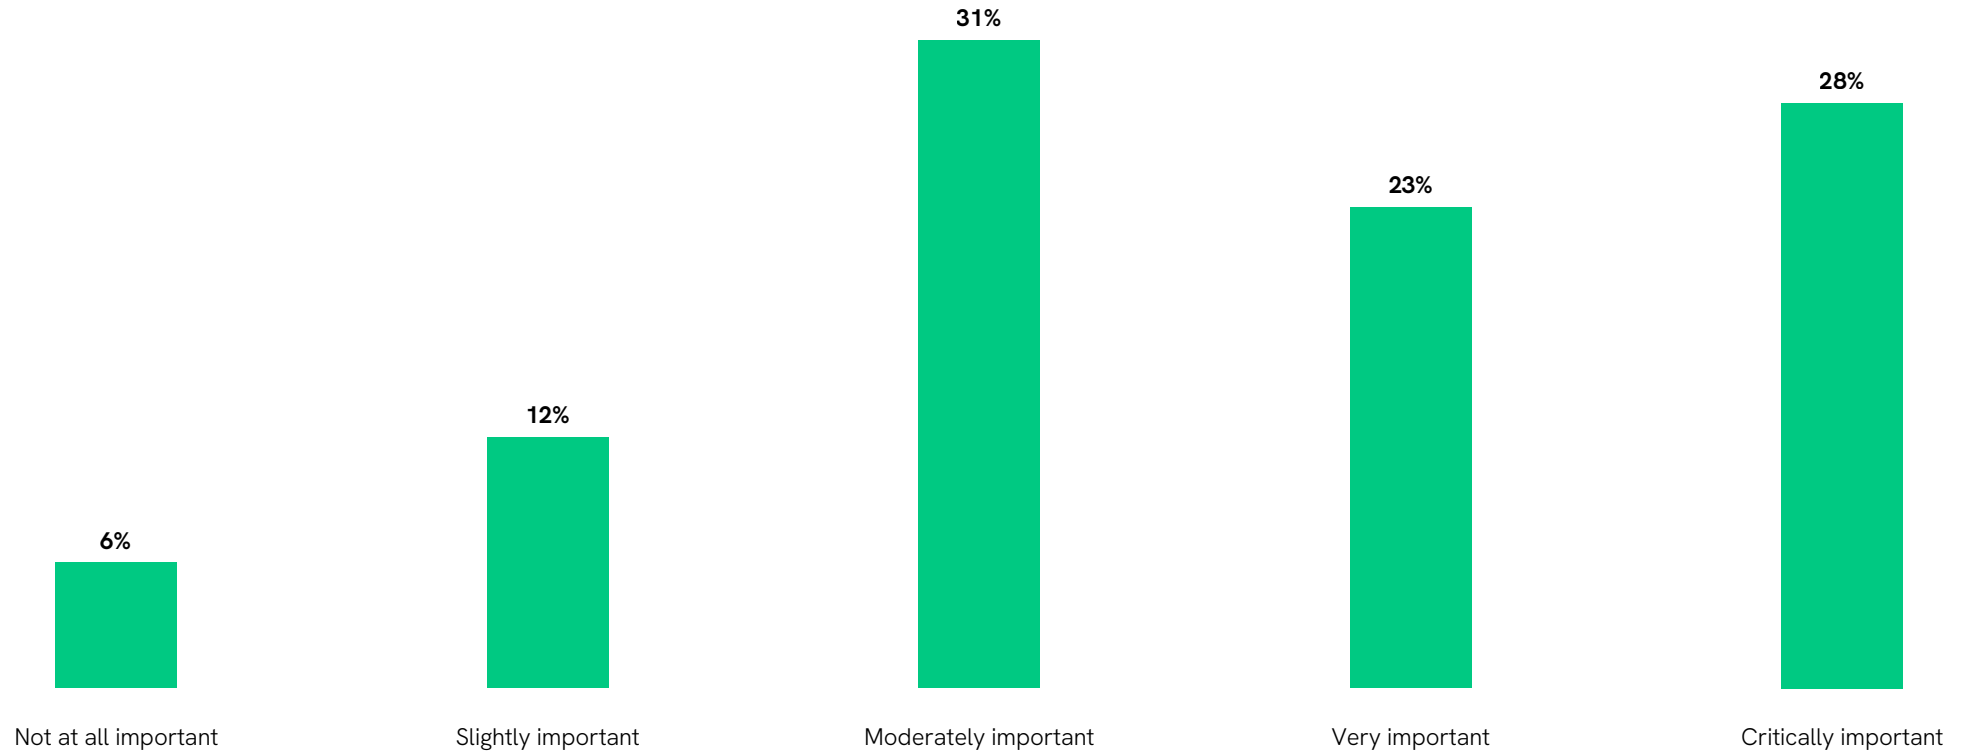
Employees' AI preparation



11. ADAPT CIO Edge Survey in Feb 2026. Sample size : 227 Australian CIOs
12. ADAPT CIO Edge Survey in Feb 2026. Sample size : 222 Australian CIOs

ADAPT's data signals add a second assurance layer. For future small language models, 51% of CDAOs say on-shore AI inference is very important or critically important. That does not make local inference a universal requirement, but it does show how quickly trust, assurance, and operating control are moving into the centre of AI decision-making.¹³

How important is on-shore AI inference (Australia-based processing) to your organisation for future Small Language Models?



13. ADAPT Data and AI Edge Survey in Apr 2026. Sample size : 156 Australian CDAOs

Platforms and partnerships are central to scaling AI capability

Once these structural constraints are visible, the platform and partner question changes. Organisations are increasingly relying on integrated platforms and trusted delivery ecosystems not simply to introduce new capability, but to reduce operating complexity, improve visibility across environments, and translate fragmented effort into more repeatable execution.

As AI and modernisation agendas mature, progress depends less on adding another disconnected layer of tooling and more on how well cloud, data, business application, security, and workflow environments work together. The next phase of value will come from improving interoperability, reducing integration friction, and extracting more from platforms already embedded in the enterprise.

This is what makes the current vendor landscape more consequential. ADAPT's research shows that enterprise environments are already relatively concentrated across a small number of providers in core business software, cloud, and data and AI platforms. In ADAPT's CIO research, Microsoft ranks as the leading primary ERP or core business software provider at 33%, ahead of Oracle and SAP. It also ranks as the leading primary cloud provider, with 67% of CIOs selecting Microsoft, ahead of AWS and GCP. In ADAPT's CDAO research, Microsoft also leads as the primary data, AI and analytics solution provider, selected by 53% of respondents¹⁴.

The significance of this is that platform position now matters more because execution increasingly depends on how easily organisations can connect data, govern usage, strengthen security, and embed AI into operational workflows without adding further complexity. Where providers already sit across multiple layers of the environment, they are often better placed to improve visibility, reduce integration overhead, and create clearer pathways from AI capability to measurable operational use.

That is why ecosystem fit is becoming a more important part of the execution question. As organisations move from experimentation to scaled deployment, they are placing greater weight on platforms and partners that can support interoperability, governance, and secure adoption across the environments they already run. In that context, vendor strength is no longer just about feature depth in one category. It is increasingly about the ability to support execution across the stack.

ERP/core business software stack provider

Rank	Parent vendor	% of respondents
1	Microsoft	33%
2	Oracle	31%
3	SAP	30%
4	Workday	5%
5	Salesforce	4%
6	TechnologyOne	4%
7	IBM	4%

Primary cloud provider

Rank	Parent vendor	% of respondents
1	Microsoft <i>(including Azure)</i>	67%
2	AWS	47%
3	GCP	5%
4	Oracle	1%

Data, AI & analytics solution provider

Rank	Parent vendor	% of respondents
1	Microsoft <i>(Including Azure, Fabric, Power BI, AI services)</i>	53%
2	Databricks	27%
3	AWS	25%
4	Snowflake	21%
5	Google	15%
6	Others	13%
7	Salesforce	9%

As the number of vendor platforms organisations rely on continues to grow, leaders are rethinking how they approach external support. Execution increasingly depends on ecosystem fit, not just technical capability in isolation. Organisations are placing greater value on delivery models that can work across existing environments, reduce complexity, and help translate platform capability into practical operating outcomes. In that context, the challenge is no longer simply choosing technology. It is orchestrating platforms, workflows, governance, and service support in ways that make scale easier rather than harder.

Security is central to that equation. As AI moves deeper into enterprise workflows, the platform decision is also a control decision. Leaders are no longer only asking whether a platform can enable AI. They are asking whether AI can be deployed into that environment with sufficient governance, identity control, auditability, data protection, and operational visibility. This is especially important as agentic and workflow-based use cases expand, because the cost of weak permissions, fragmented oversight, or poorly governed integrations rises quickly once AI moves closer to operational decision-making. In practice, the question is not just whether the environment can support AI technically, but whether it can support it securely, responsibly, and at scale.

As execution becomes harder to de-risk, the criteria for external support are becoming more selective. Senior leaders devote limited time to new vendors, so relevance and credibility matter disproportionately. Among CIOs, 54% say it is very important or essential to meet a vendor’s local executive leadership before committing to a major purchase. Among CFOs, that rises to 64%. Local proof matters too. In ADAPT’s 2026 CIO survey, 58% of CIOs say localised success stories are most useful when engaging prospective vendors, and 46% cite the importance of ROI and business case modelling support¹⁵. The signal is not simply that buyers prefer local engagement. It is that partner and ecosystem decisions are increasingly filtered through trust, contextual relevance, and confidence in execution.

When meeting with prospective vendors, what type of information is most useful to you at the early (pre-sale) stage?



15. ADAPT CIO Edge Survey in Feb 2026. N = 215 Australian CIOs

The market is therefore moving away from asking who can help launch an initiative toward asking which combination of platforms, expertise, and delivery support can help operationalise, govern, and sustain value at scale. That is a more demanding test. It places greater weight on interoperability, accountability, local executive access, and the ability to turn broad platform capability into outcomes that fit the organisation’s actual operating environment.

Shift towards managed services and partner-led execution

This shift also helps explain the growing relevance of managed services and partner-led execution models. As organisations confront uneven data maturity, constrained internal capacity, governance pressure, and rising expectations around value realisation, the challenge is no longer just delivery. It is sustained execution.

In that environment, managed services become more than an outsourcing model. They become a way to reduce the execution gap. Where internal teams are stretched across legacy modernisation, governance uplift, security requirements, and AI experimentation, partner-led models can help organisations accelerate progress by bringing operating discipline, specialist capability, and continuity across the environments that matter most.

This is particularly relevant where AI depends on several conditions being improved at once: cleaner data foundations, stronger governance, better integration across cloud and business applications, tighter security controls, and more consistent day-to-day operational support. Very few organisations can mature all of those layers at speed through internal effort alone. Managed services can therefore play a useful role when they help simplify delivery, improve operational resilience, and provide a clearer path from experimentation to repeatable execution.


The value of this model is not simply cost efficiency. It is execution confidence. The stronger managed-service proposition is the ability to help organisations sustain and govern AI-enabled environments over time, not just deploy them. That includes maintaining platform health, supporting secure integration, strengthening visibility across workflows and users, and keeping environments aligned to changing governance and business requirements.

This is where the market is becoming more selective. Partner-led execution is most valuable when it reduces fragmentation, clarifies accountability, and helps organisations move faster without weakening control. In practice, the most relevant providers will be those that can combine ecosystem familiarity, security discipline, local delivery credibility, and managed operational support into a model that helps clients close the gap between ambition and execution.

Sector operating realities

A useful starting point is the broader execution baseline across the market. On average, only 61% of current technology capabilities are adopted and in use, indicating that a significant share of technology investment is still not translating into practical enterprise value. That matters because AI is not being introduced into clean environments. It is being layered into organisations already managing underused platforms, fragmented workflows, and uneven operating maturity. In this context, the more relevant question is not simply which sectors are investing in AI, but which sectors are structurally better positioned to move from pilots to repeatable outcomes.

Adoption and use of current technology capabilities - by sector (including tools, platforms, and applications currently in use across the business)



Three benchmark lenses help make that clearer

1

First, data preparedness to support AI at scale shows a market that is still largely in transition. Healthcare appears relatively stronger, with 43% in the developing stage and no respondents yet describing themselves as optimised or fully ready. Financial services is slightly more advanced, with 35% developing, 11% optimised, and a small fully ready cohort. Government remains more constrained, with 27% not ready and 33% emerging, meaning six in ten agencies are still below a genuinely scalable data footing. Education shows a similar pattern, with 15% not ready and 62% still emerging, suggesting that foundational work remains the dominant task.

2

Second, architecture maturity reinforces that most sectors are not yet able to scale GenAI without further structural change. Financial services looks comparatively stronger here, with 19% highly capable and the rest largely split between limited support and partial capability. Healthcare is somewhat further forward than government and education, with 21% highly capable and 4% fully future-ready. Government remains constrained, with 13% saying their architecture is not at all ready, 47% saying support is limited, and none yet describing it as highly capable or future-ready. Education is similarly early, with 17% not ready, 50% limited support, and only 8% highly capable.


3

Third, measurable ROI from GenAI shows that most sectors are still proving value rather than scaling it. Financial services sits in the middle: 22% report no measurable ROI to date, 47% see only early signs, and 24% say ROI is broadly meeting expectations, with very little above that. Healthcare is somewhat stronger on broad value realisation, with 32% saying ROI is meeting expectations, but only 5% say it is exceeding expectations in some areas. Government remains earlier, with 20% reporting no measurable ROI and 70% seeing only early, limited, or unclear returns. Education is similar, with 25% reporting no measurable ROI, 58% seeing only early signs, and none reporting ROI exceeding expectations.


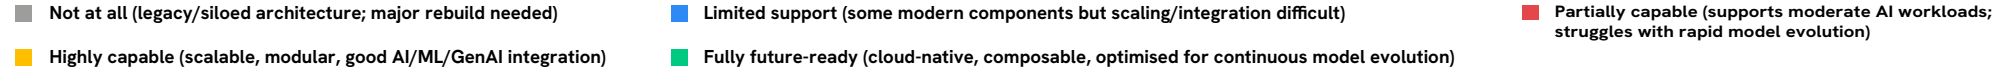
Taken together, these benchmarks point to a clear pattern: financial services and healthcare are somewhat further ahead on the structural conditions for AI scale, while government and education remain more weighted toward foundational work and early value realisation. That benchmark backdrop helps explain why the same execution challenge plays out differently by sector.

Data preparedness to support AI at scale


- 1 — Not ready (significant data gaps; inconsistent, siloed, or low-quality data)
- 2 — Emerging (foundational data work underway; partial quality and governance in place)
- 3 — Developing (scalable data pipelines in progress; moderate consistency and accessibility)
- 4 — Optimised (high-quality, well-governed, reliable data supporting advanced use cases)
- 5 — Fully ready (enterprise-grade, AI-optimised data foundation enabling scale across functions)



Data & AI architecture maturity to enable genai without major re-engineering



Measurable ROI from GenAI investments



Financial services remains one of the clearest examples of AI ambition running into structural execution limits. Relative to several other sectors, it is better positioned on data preparedness and architecture maturity, but not yet to the point where broad, lower-friction scale is assured. The sector's main constraint is not tool access. It is the difficulty of moving data with enough consistency, speed, and control across complex, high-liability environments.

That is why AI is scaling most credibly in bounded, defensible domains such as process automation, risk detection, compliance support, and decision support for operations, finance, and fraud. These are use cases where data is relatively better defined, governance expectations are clearer, and value can be demonstrated without introducing broad autonomy into core decision environments. The benchmark data supports that interpretation: financial services shows some movement into optimised data readiness and highly capable architecture, but ROI still clusters heavily in the "early signs" and "meeting expectations" range rather than in stronger return categories. The sector is therefore best described as structurally ahead of much of the market, but still progressing through controlled, measured scale rather than broad operational breakthrough.

Healthcare

Healthcare shows a different pattern. Among the priority sectors, it is one of the better-positioned environments on data preparedness, and it also shows relatively stronger architectural support than government or education. But the issue here is not simply readiness. It is the ability to translate relatively stronger trust and governance conditions into timely, context-rich operational use.

That is why the most credible momentum is likely to remain in areas such as administrative productivity, knowledge management, operational support, cyber threat detection, and service experience rather than highly autonomous or deeply orchestrated use cases. The benchmarks support that view. Healthcare has fewer organisations in the "not ready" category than government or education, and a larger share saying ROI is broadly meeting expectations. But it still shows only a small minority with stronger return profiles, and only a limited share with future-ready architecture. In other words, healthcare appears better placed to operationalise AI in contained, high-value domains, but not yet uniformly ready to scale it across more complex, real-time care and operational environments.

Government reflects the same execution logic under the added weight of public accountability, cross-agency coordination, and service continuity. The benchmark data makes clear that this sector remains more heavily concentrated in foundational work than the others. On data preparedness, 60% of agencies are still in the not ready or emerging categories. On architecture, 60% sit in not ready or limited support, and none yet describe their environment as highly capable or future-ready. On ROI, the picture is even earlier: 90% of agencies report either no measurable ROI or only early signs with limited or unclear returns.

This helps explain why AI in government is likely to scale first in contained, lower-liability domains such as fraud detection, cyber threat detection, knowledge management, administrative automation, and code testing. These are use cases that can sit within clearer data boundaries and stronger accountability settings. What remains difficult is broader orchestration across agencies, more real-time multi-system decisioning, and more autonomous public-facing applications. In government, the constraint is not a lack of interest. It is the difficulty of combining interoperability, assurance, and speed in the same operating environment.

Higher education

Higher education sits closer to government than to financial services or healthcare in its current benchmark profile. It remains heavily weighted toward foundational work on both data preparedness and architecture, with 77% still in the not ready or emerging categories on data and 67% in not ready or limited-support architecture environments. ROI is similarly early: 83% report either no measurable value or only early, limited, or unclear returns.

That suggests higher education is still earlier in the journey from experimentation to repeatable institutional value. The most credible near-term pathways are therefore likely to centre on administrative efficiency, staff productivity, knowledge access, service support, and selected teaching or research enablement use cases rather than broad autonomous deployment. The challenge in higher education is less a lack of AI activity than the difficulty of aligning governance, data access, architecture, and operating capability across teaching, research, and student-service environments. In that sense, the next question for the sector is not how much experimentation is happening, but how much of it can be sustained and scaled without adding disproportionate complexity.

AI as a capacity multiplier, not a frontier play

Not-for-profit organisations are approaching AI less as a frontier innovation play and more as a practical way to extend limited organisational capacity. The key question in this segment is not how quickly advanced AI can scale, but where it can reduce administrative burden, strengthen service responsiveness, improve knowledge access, and help lean teams operate more efficiently without adding complexity they cannot absorb.

The maturity profile remains early. Most organisations are still at the awareness or exploration stage for agentic AI, with very few progressing into targeted implementation and none yet reporting broader operational scaling or value optimisation. The main barriers are also foundational rather than strategic: data readiness is the single largest constraint, followed by workforce readiness. This suggests the limiting factor is not appetite for AI, but whether the underlying environment is strong enough to support broader use safely and effectively.

Governance and workforce preparation are developing, but remain lightweight. Most organisations have usage guidelines or draft principles in place, and many are encouraging AI-assisted workflows with optional upskilling. Far fewer have more formal structures such as dedicated AI champions, centres of excellence, or enterprise-wide governance with strong board involvement. This points to a segment that is trying to build safe and practical AI foundations, but doing so with lean operating models and limited specialist capacity.

Use-case adoption is concentrated in areas that align with immediate operational value. The strongest traction is in operational improvements, common administrative tasks, cyber security threat detection, customer service, and knowledge-heavy work. That pattern suggests AI is being used primarily to support day-to-day efficiency and staff productivity, rather than to drive broad autonomous decision-making.

There is no direct benchmark showing that AI is already improving funding efficiency or solving resourcing pressure at scale across the segment. However, the use-case pattern does indicate that organisations are treating AI as a capacity extender: focusing on assistive, governed, and relatively contained applications that can help constrained teams do more with limited time and capability. In this segment, success is likely to come less from the breadth of deployment and more from the quality of fit. Simplicity, governance, and absorptive capacity matter as much as technical capability.

What this means for leaders

The leadership challenge is no longer to prove that AI matters. It is to decide where it should scale, under what conditions, and with what level of organisational readiness. The market has already moved beyond broad experimentation. The next phase will be defined by which organisations can convert AI ambition into practical, governed, and repeatable execution. That requires a more disciplined leadership agenda across five areas.

1. Scale what is credible, not just what is exciting

The strongest organisations will not be those pursuing the widest spread of AI activity. They will be those making clearer decisions about where AI can create measurable value, where workflows are stable enough to support change, and where outcomes can be governed with confidence. Selective scale is likely to prove more valuable than broad but uneven rollout.

2. Treat data and operating maturity as strategic conditions

The real constraint is increasingly not access to AI capability, but the strength of the environment beneath it. Leaders should treat data readiness, integration depth, architecture flexibility, governance coverage, and operating discipline as strategic conditions for scale. Where those remain weak, AI will stay bounded regardless of executive ambition.

3. Modernise with a view to flow, not just platform progress


Modernisation only creates strategic advantage when it reduces friction across the organisation. That means simplifying handoffs, improving data movement, strengthening visibility across systems, and enabling more consistent delivery and decision-making. Platform change on its own is not enough. The real test is whether it improves the organisation's ability to operate, adapt, and scale with less drag.

4. Strengthen accountability before increasing autonomy

As AI becomes more embedded in workflows and decision environments, leadership responsibility becomes sharper, not lighter. Governance ownership, escalation paths, risk oversight, and business accountability need to become clearer as capability becomes more distributed. The more organisations pursue automation and agentic use cases, the less room there is for ambiguity over who owns the decision, the control, and the outcome.

5. Lead through ecosystems, not isolated technology decisions

A growing share of execution success now depends on how well organisations align the platforms, partners, and internal capabilities already around them. Most enterprises are not building AI in isolation. They are working across existing cloud, data, productivity, business application, and workflow environments, often with multiple internal and external contributors. That means the leadership task is increasingly one of orchestration: ensuring that ecosystem choices reduce fragmentation, improve interoperability, and create clearer pathways from capability to value.



A practical framework for leaders: five tests of scalable execution

Leaders should increasingly evaluate major AI and modernisation decisions against five questions:

1

Is the use case worth scaling?

Can it create measurable value beyond technical success?

2

Is governance strong enough?

Can the organisation monitor, explain, and defend what is being deployed? Will this reduce or increase friction?

3

Does the ecosystem strengthen execution?

Will the chosen mix of platforms, partners, and internal ownership make value easier to operationalise over time?

4

Is the enterprise ready enough?

Are the data, workflow, architecture, and operating conditions strong enough to support scale? Is governance strong enough?

5

Will this reduce or increase friction?

Does the decision simplify the environment, or add another disconnected layer? Does the ecosystem strengthen execution?

The next phase of advantage will come from better sequencing. The organisations that progress fastest will not be those that pursue the most AI activity, but those that are clearest about what is ready to scale, what still requires foundational work, and what should remain contained until the operating environment is stronger.

Recommendations: Closing the execution gap and accelerating AI outcomes

The next phase of AI progress will depend less on ambition and more on the ability to close the gap between experimentation and repeatable execution. Six priorities stand out.

1. Strengthen the foundations first

AI scale still depends on better data quality, accessibility, governance, and interoperability. Where foundations remain weak, outcomes will stay fragmented.

2. Tie AI to measurable outcomes

Investment decisions need clearer links to productivity, service improvement, risk reduction, or business value. The priority is not more activity, but stronger paths to impact.

3. Prioritise use cases that reduce friction

The most credible early gains are likely to come from practical applications that simplify workflows, support decision-making, and help constrained teams operate more effectively.

4. Build security and governance into scale plans

AI will not scale credibly without stronger controls around access, oversight, auditability, and secure deployment. Governance and security need to move closer to delivery.

5. Use ecosystems and managed execution to accelerate progress

The right mix of platforms, partners, and managed support can help reduce complexity, improve interoperability, and sustain AI-enabled environments over time.

6. Sequence effort where readiness is strongest

Progress will remain uneven across sectors and functions. The strongest results will come from focusing first where data, architecture, governance, and workflow conditions are already more mature.

What stronger execution requires in practice

For many organisations, closing the execution gap will require more than new tools or isolated pilots. It will require a clearer line between data, governance, security, platform integration, and day-to-day operational support. The organisations moving most effectively are not simply experimenting more. They are building trusted data foundations, improving performance visibility, reducing delivery complexity, and creating more secure pathways from AI ambition to measurable outcomes.

In practice, that means strengthening the conditions that make AI useful at scale: reliable and governed data, clearer decision support, better integration across cloud and business environments, stronger control over risk, and the operational capability to sustain progress over time. This is where execution support becomes most valuable, not as an overlay to innovation, but as part of how organisations turn strategic intent into repeatable business impact.

Conclusion

Australia’s enterprise market is not short on AI ambition. What is becoming clearer is that ambition alone is no longer a differentiator. The organisations that will make the most progress over the next phase are those that can turn ambition into practical execution without increasing complexity faster than they create value.

That is why execution has become the defining lens. AI outcomes are now being shaped by the strength of data foundations, the maturity of modernisation efforts, the depth of governance and assurance, and the organisation’s ability to work effectively across the platforms, partners, and operating environments already in place.

The strategic question, then, is no longer whether AI belongs on the agenda. It is whether the enterprise can support it in a way that is measurable, governed, secure, and repeatable. In that environment, success will come less from pursuing the most activity and more from making better decisions about where to scale, how to reduce friction, and which ecosystems can help turn capability into durable business outcomes.

Guide to market segmentation used across surveys:

Large enterprise	(>10,000 emp.)
Enterprise	(2501 to 10,000 emp.)
Mid-market	(501 to 2500 emp.)
Small / medium corporate	(1-500 emp.)

Nexon Asia Pacific


Nexon Asia Pacific (Nexon) is an award-winning digital & IT services partner for mid-market, enterprise & government organisations across Australia. We have a uniquely broad suite of solutions to service clients who require end-to-end capabilities coupled with specialist expertise in security, cloud, and digital solutions.


As a certified and accredited local and state government provider, CREST- and ISO-certified, Nexon partners with world-class technology vendors to deliver innovative and integrated solutions.

Pointedly structured and culturally driven to deliver the highest standards of responsiveness, competency, and transparency, our integrated solutions and close client relationships help propel organisations forward. For expertise delivered with care, clients can look to the future with Nexon.




Take the next step.
Get in touch with the Nexon team.

 1300 800 000

 nixon.com.au

ADAPT

 +61 (2) 9435 3535

 hello@adapt.com.au

 adapt.com.au

ADAPT is a specialist Research & Advisory firm providing local market insights and benchmarking data to Australia and New Zealand's senior technology and business community.

Since 2011, through our proprietary surveys, we have empowered the region's top enterprise and government leaders to make informed strategic decisions and extract greater value from their technology investments.

Combined with our industry-leading conferences, private roundtables, and custom research projects, ADAPT equips organisations with the evidence, connections, and tools they need to stay at the forefront of modern trends and build for the future.

Disclaimer: ADAPT's research reports are based on proprietary survey insights and expert analyst opinions, curated to inform strategic decision-making. The report findings represent independent analysis and should not be considered definitive statements of fact. ADAPT does not endorse any vendor, product, or service. All research is provided "as is" without warranties, including those of merchantability or fitness for a particular purpose.