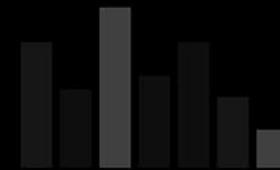


ADAPT

# Emerging Technology Trends

Australia & New Zealand 2026



# Executive summary

Australian organisations are entering a decisive phase in their technology transformation journey. After several years of experimentation with cloud, digital platforms, and generative AI, the market has reached a point where execution discipline, sequencing, and foundations now matter more than ambition.

## Current challenges



**76%**

of organisations are **NOT AI-data ready**.



**41%**

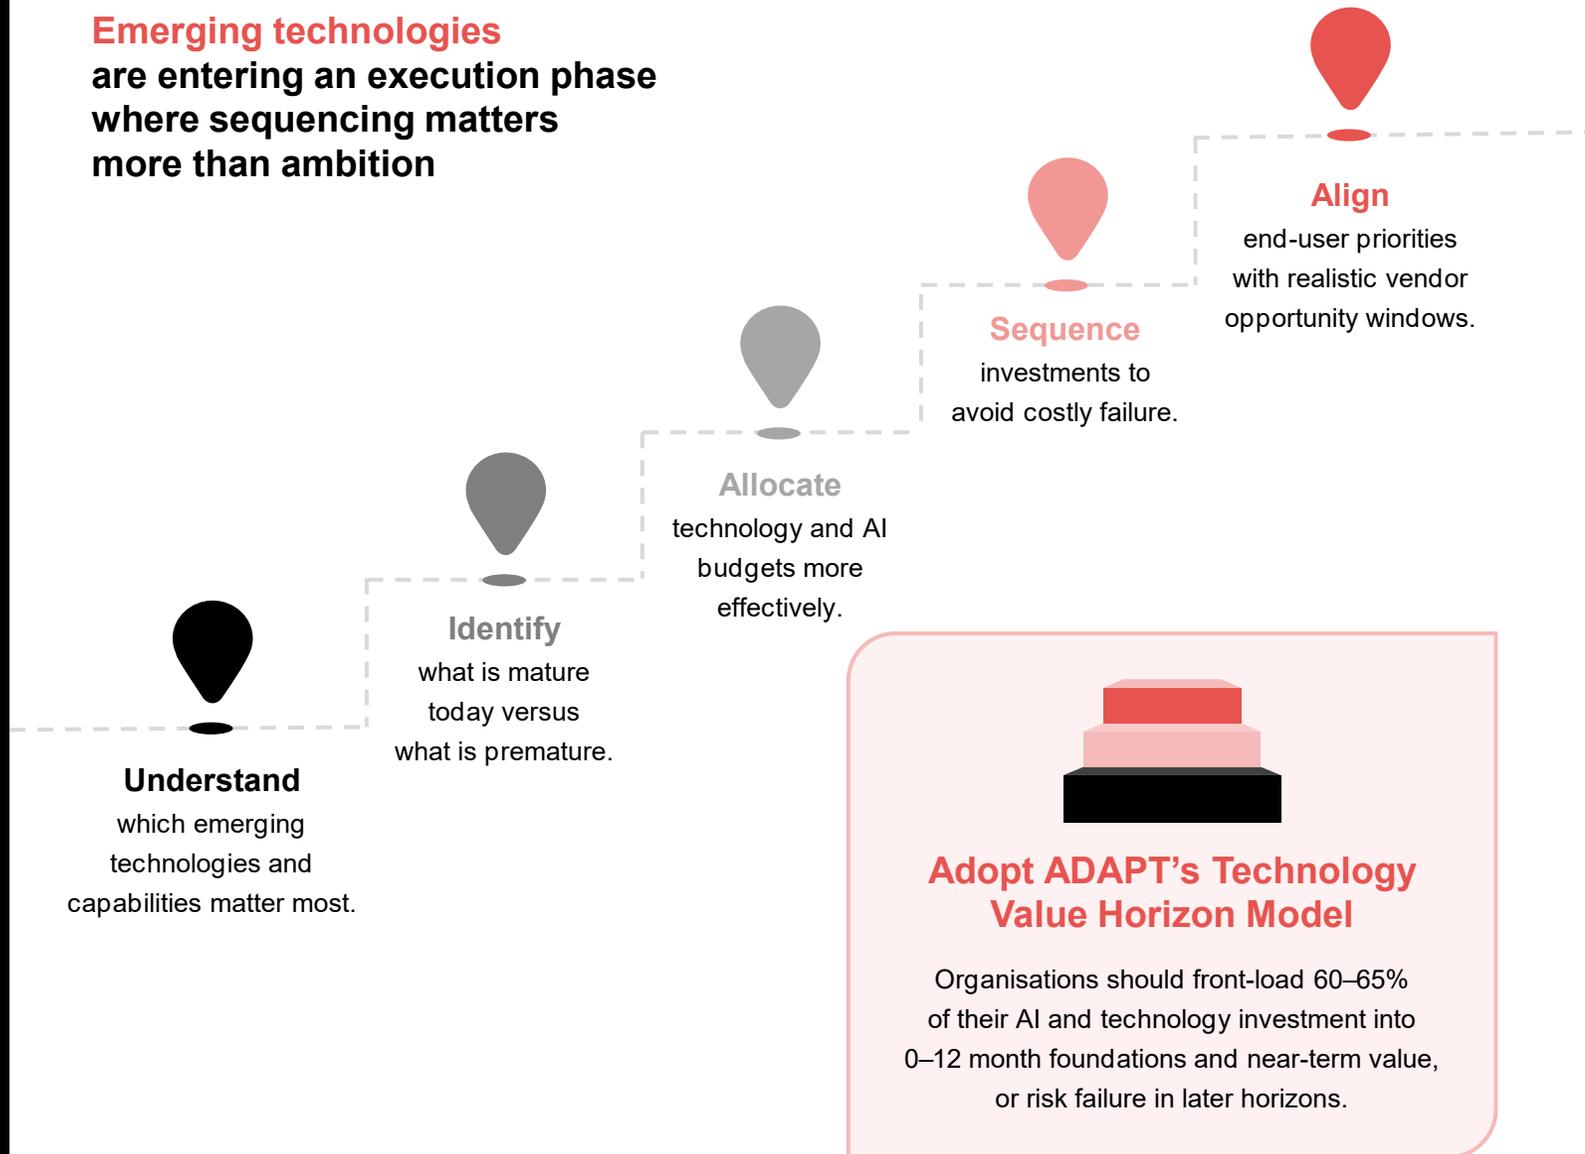
of mission-critical apps run on **legacy platforms**.



**25%**

of enterprises are **repatriating cloud workloads**.

**Emerging technologies** are entering an execution phase where sequencing matters more than ambition



# Research methodology

**900+** ADAPT surveys and interviews with **900+ Australian CIOs, CISOs, CFOs, CDAOs, and infrastructure leaders.**

Analysis focused on maturity, investment intent, dependencies, and ROI timelines.



Maturity



Investment intent



Dependencies



ROI timelines

Coverage across enterprise, mid-market, and public sector.



Enterprise



Mid-market



Public sector

## Analysis

This report provides a disciplined, horizon-based framework to help CIOs prioritise investments, reduce execution risk and allocate capital effectively over the next three years.

03

### Horizon 3 (3+ years): **Selective autonomy**

Real-time architectures, agentic systems, and next-generation infrastructure become viable only after foundational and | operational capabilities are established.

02

### Horizon 2 (18-24 months): **Scale**

Embed AI into workflows and modernise core platforms to enable reuse, integration, and operational efficiency.

01

### Horizon 1 (0-12 months): **Foundations**

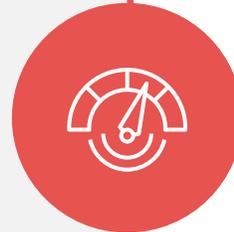
Focus on data quality, security, cost control, and core platforms. Determine whether AI and automation can scale safely or deliver ROI. Skipping this stage consistently leads to failure in later horizons.

# Report overview

Australian organisations are moving from AI and cloud experimentation to execution, but gaps in readiness, sequencing, and operating discipline remain the main obstacles.

This report highlights the technologies and capabilities that matter most in the next three years, clarifying what to invest in now versus later.

Using ADAPT's 2025 research, it provides practical guidance on what to prioritise, scale, or defer based on maturity, dependencies, and real operational impact.



## How to use this report

### **Prioritise investments**

Use the horizon framework to distinguish foundational investments from scaling and longer-term bets and allocate capital accordingly.

### **Sequence initiatives to avoid premature AI failure**

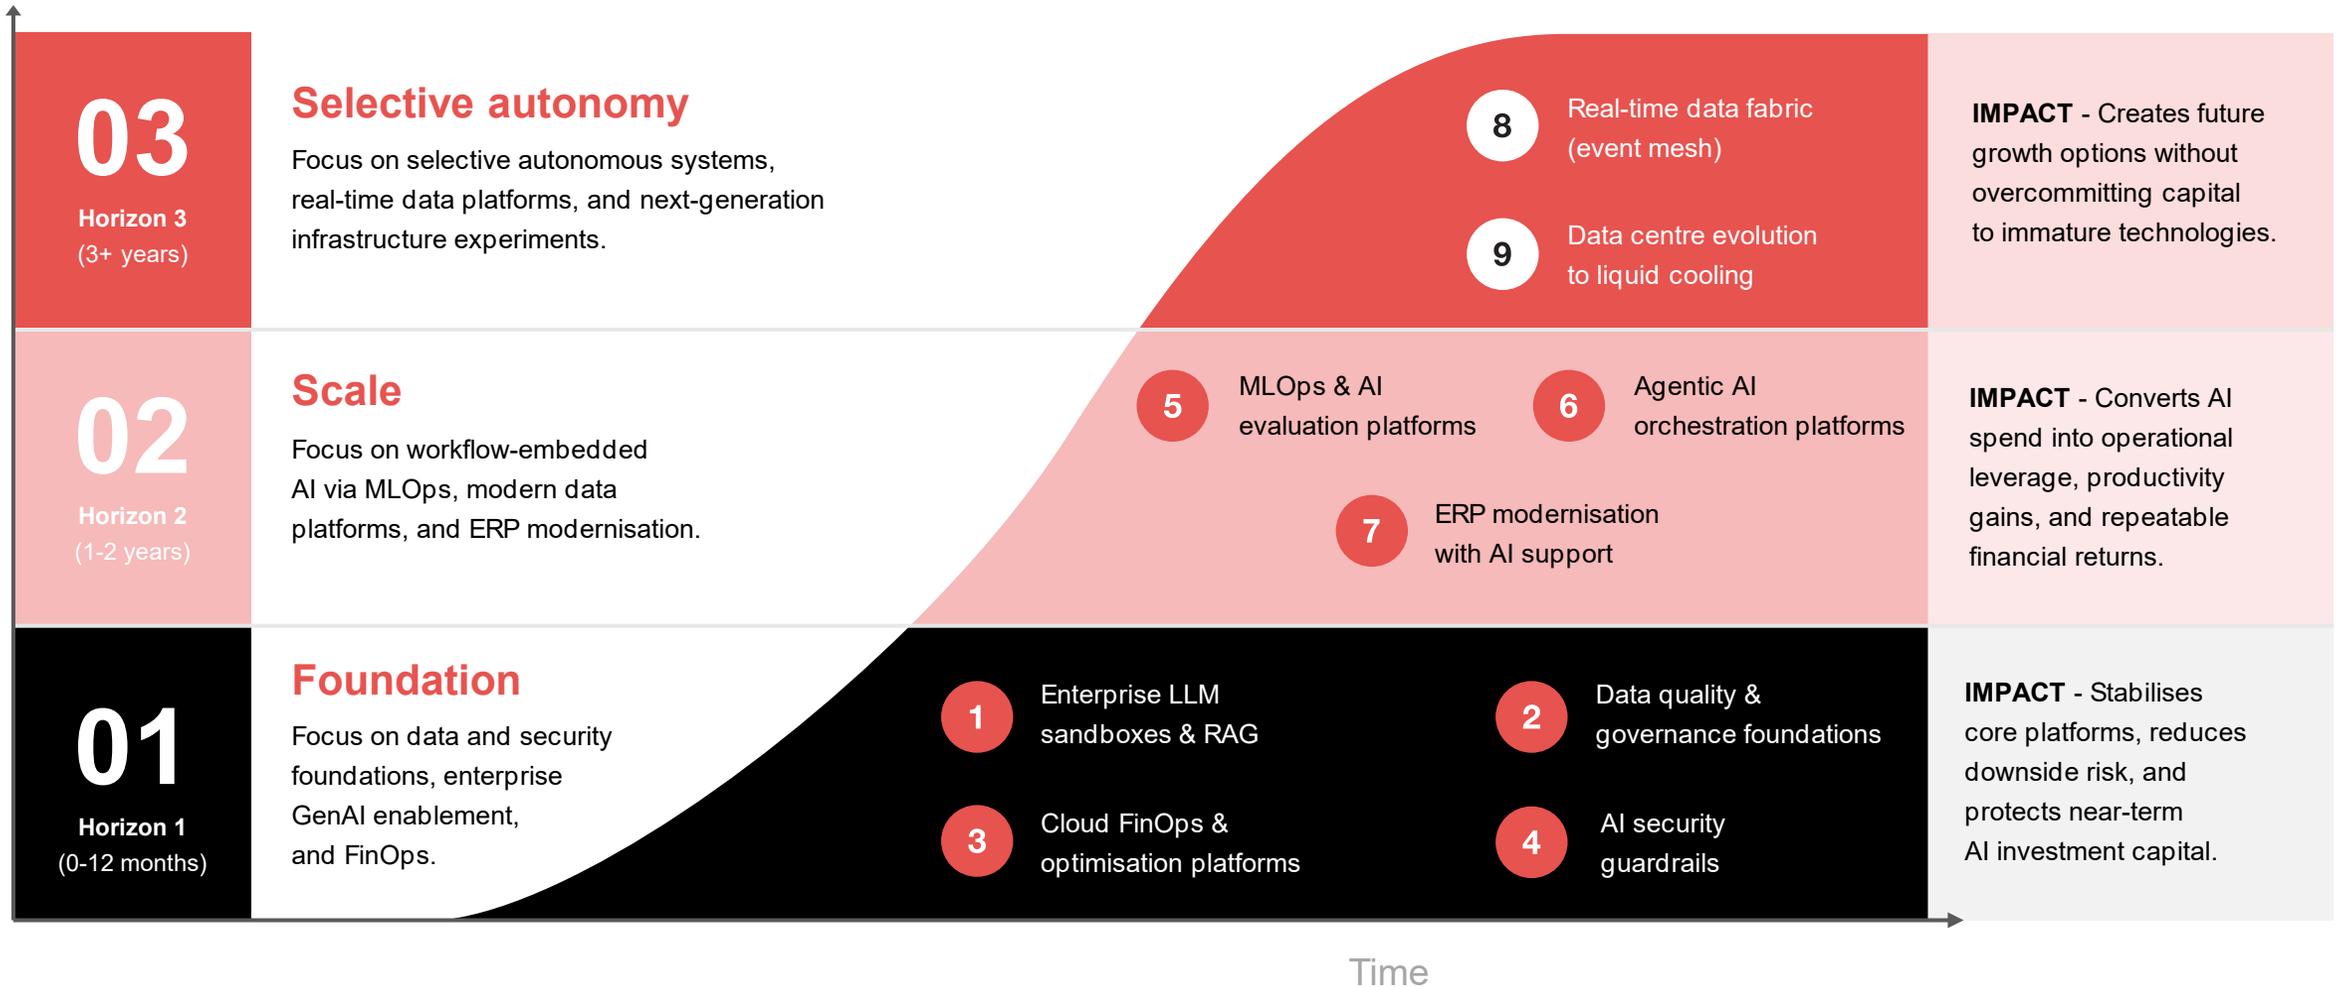
Apply dependency logic to ensure data, security, and operating foundations are in place before scaling AI and automation.

### **Align executives on realistic success metrics**

Use the CIO-level KPIs to set shared expectations on adoption, control, and time-to-value across leadership teams.

# ADAPT Emerging Tech Trends 2026

Value and operational complexity

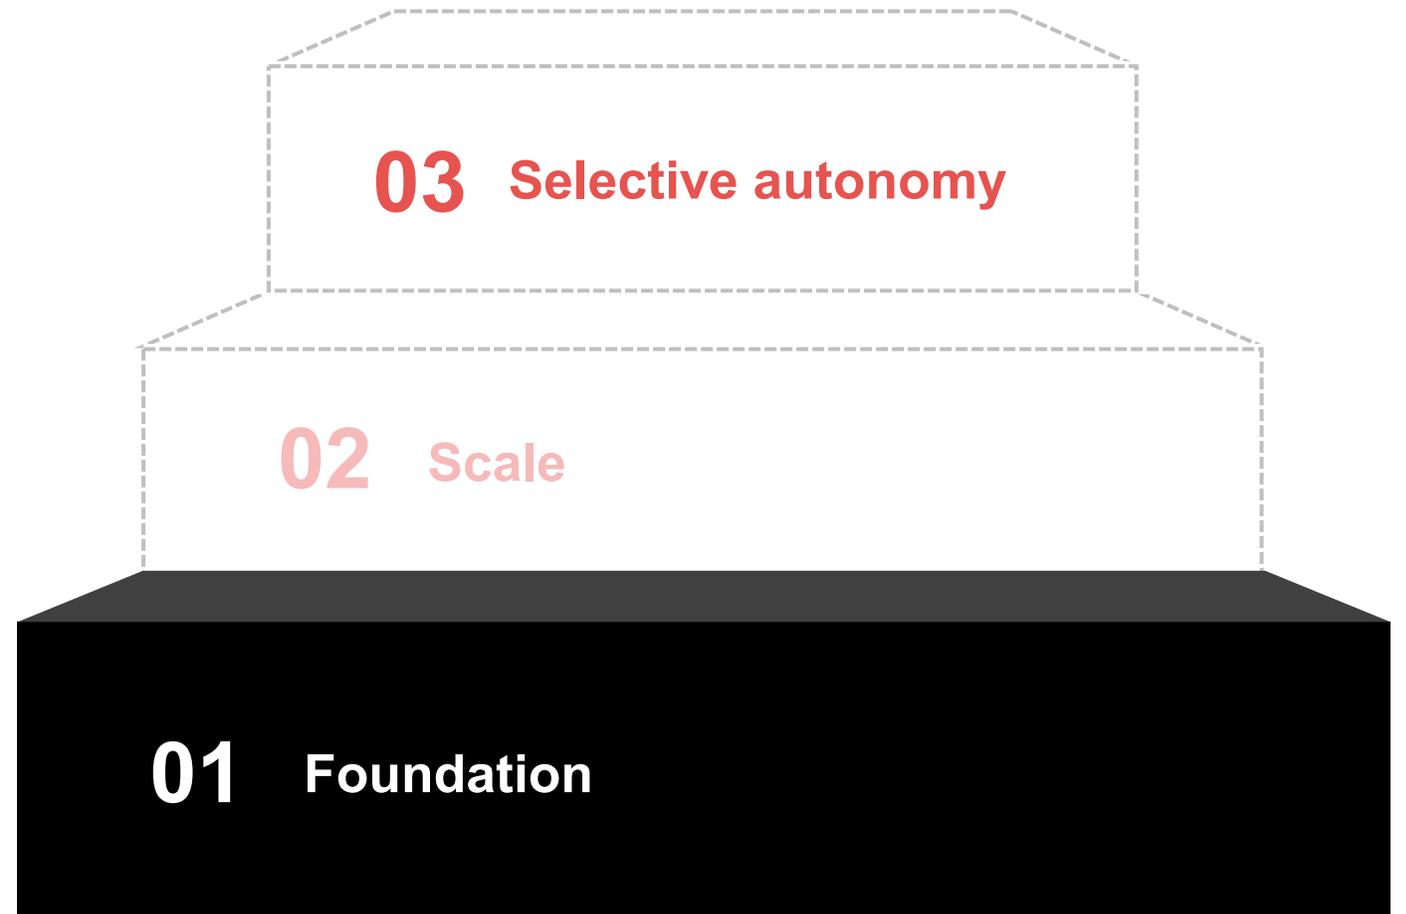


# Deep dive: Emerging tech trends



# 01

## Foundation



# 01 Foundation

## Trend 1

# Enterprise LLM sandboxes & RAG

<b>Role</b>	Foundation for immediate GenAI value
<b>Definition</b>	A <b>secure enterprise platform</b> that allows employees to use GenAI on <b>internal, governed data</b> without exposing sensitive information or creating security risk.
<b>Why this matters now</b>	This is the most mature GenAI capability with <6-month ROI and is a required dependency for all downstream GenAI and AI-enabled workflow initiatives.
<b>Capabilities</b>	<ul style="list-style-type: none"> <li> Knowledge retrieval and summarisation across operations, legal, HR, and IT.</li> <li> Agent-assist for call centres and service teams.</li> <li> Research and content support for internal staff.</li> <li> Code assistance and documentation for engineering teams.</li> </ul>
<b>Investment profile</b>	<ul style="list-style-type: none"> <li>• <b>Budget:</b> ~25–30% of total AI / tech spend.</li> <li>• <b>Timeline:</b> 3–6 months for initial deployment and value realisation.</li> </ul>



## Success metrics



### Adoption and usage

- **30–50%** of eligible knowledge workers active within 90 days.
- Weekly active users / eligible users **≥35–45%**.
- **≥3–5** GenAI sessions per active user per week.



### Productivity & efficiency

- **20–30%** reduction in time spent on search, drafting, and summarisation tasks.
- **10–20%** handle-time reduction in service and contact-centre workflows.
- Document / research turnaround time reduced by **≥25%**.



### Risk & control

- Zero critical data leakage or policy violations.
- **100%** of GenAI activity routed through approved, logged environments.
- **≥95%** adherence to prompt, data access, and usage policies.



### Platform readiness

- RAG reused across 3+ functions.
- Standardised LLM access layer supporting multiple models/vendors.
- Clear integration path into copilots, MLOps.

## Market signals



Ask Telstra is an internal generative AI assistant enabling employees to search and retrieve information from enterprise knowledge [bases](#).



Australian Government

GovAI provides APS agencies with a secure, centralised GenAI sandbox for controlled experimentation and [deployment](#).

ANZ is advancing enterprise GenAI adoption through an AI Immersion Centre and workforce-wide Copilot [enablement](#).

## Sample vendors

### Global

- **Microsoft:** Azure OpenAI Service
- **AWS:** Amazon Bedrock (Knowledge Bases)
- **Google Cloud:** Vertex AI Search / Agent Builder

### Local

- **Blue Crystal Solutions:** private LLM Service
- **Telstra Purple:** enterprise GenAI Platform (managed)
- **Vault Cloud:** sovereign GenAI environment



Applied safely, generative AI can enhance experiences for customers and employees, and help us generate business outcomes faster, with lower risk and at lower cost."

### Gerard Florian

ANZ Group Executive of Technology

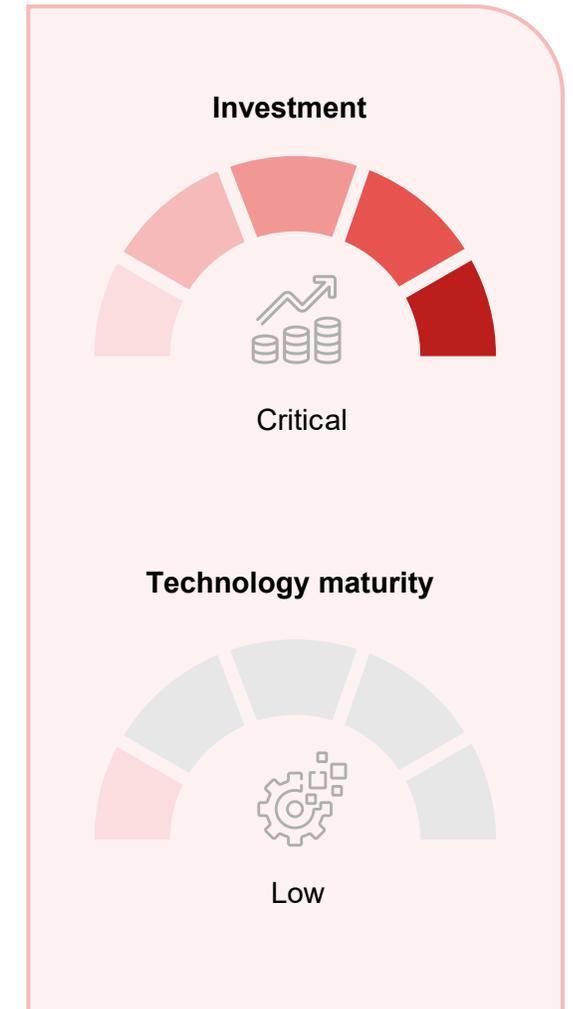
Read more 

# 01 Foundation

## Trend 2

# Data quality & governance foundations

<b>Role</b>	Critical enabler for scalable and trusted AI
<b>Definition</b>	An enterprise capability that ensures data used across analytics and AI is accurate, consistent, traceable, and governed, with clear ownership and controls.
<b>Why this matters now</b>	Most AI initiatives stall or fail due to poor data quality and unclear governance. Strengthening data foundations is a prerequisite for scaling GenAI, analytics, and automation with confidence.
<b>Capabilities</b>	<ul style="list-style-type: none"> <li> Standardised enterprise data definitions with clear ownership and accountability.</li> <li> Data quality controls embedded into critical data pipelines and reporting flows.</li> <li> End-to-end data lineage and traceability across analytics, AI, and regulatory use.</li> <li> Policy-based data access and governance aligned to risk, privacy, and compliance requirements.</li> </ul>
<b>Investment profile</b>	<ul style="list-style-type: none"> <li>• <b>Budget:</b> ~20–25% of total AI / tech spend.</li> <li>• <b>Timeline:</b> 9–12 months for foundation uplift and measurable impact.</li> </ul>



## Success metrics



### Ownership & adoption

- **70–90%** critical domains owned.
- **60–80%** priority datasets certified.
- **50–70%** analytics via governed layers.



### Data quality outcomes

- **20–40%** reduction in data rework.
- **25–50%** fewer quality-related defects.
- **15–30%** faster analytics delivery.



### Risk, compliance & control

- **90–100%** lineage for critical datasets.
- **80–95%** fewer manual access exception.



### AI enablement readiness

- **60–80%** AI-relevant datasets qualified.
- 3–5 reusable data products live.
- Clear path to MLOps & GenAI integration.

## Market signals



### Australian Taxation Office (ATO)

Enterprise-wide data governance and stewardship supporting trusted analytics and regulatory [reporting](#).



### NSW Government

State-wide data governance and quality standards enabling cross-agency analytics and [decision-making](#).

## Sample vendors

### Global

- **Collibra:** data intelligence cloud
- **Informatica:** intelligent data management cloud, data governance & quality
- **Alation:** data catalogue & governance platform

### Local

- **Precisely:** data integrity suite
- **Vault Cloud:** sovereign data platforms, governance-ready environments
- **Databricks ANZ:** unity catalogue



Effective governance balances worst-case threat preparedness with the need to drive value from new technology adoption."

### Alex Serrano

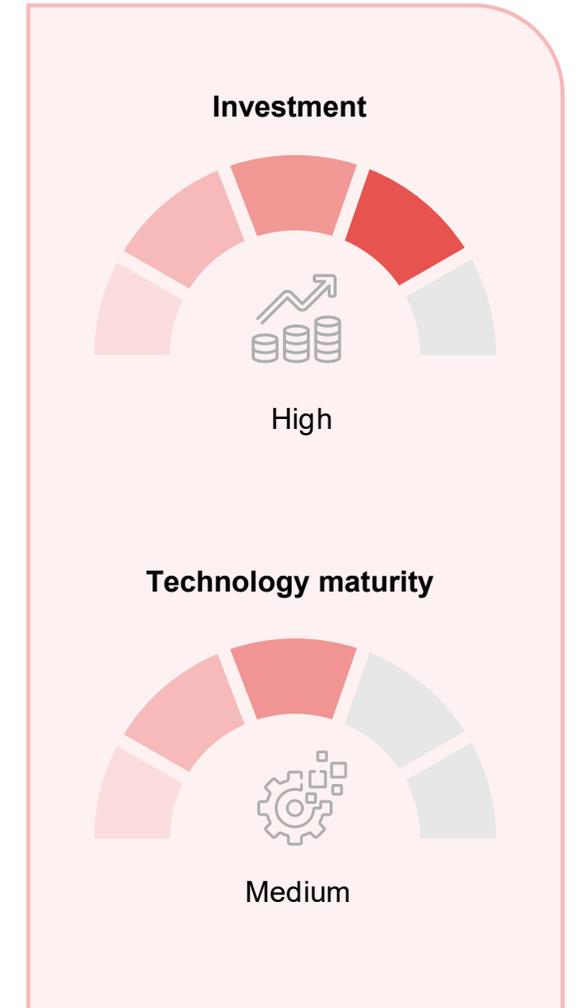
Technology Governance  
Lead at SEC Victoria

Read more [————>](#)

## Trend 3

# Cloud FinOps & optimisation platforms

<b>Role</b>	Restoring cost discipline and capital efficiency in cloud spend
<b>Definition</b>	An enterprise capability that provides visibility, allocation, and control of cloud spend, enabling technology and finance leaders to actively manage costs while supporting business demand.
<b>Why this matters now</b>	Cloud costs have become structural and opaque, driven by overprovisioning, fragmented ownership, and GenAI workloads. FinOps enables CIOs to reassert financial control without slowing innovation.
<b>Capabilities</b>	<ul style="list-style-type: none"> <li> Real-time cloud cost visibility by application, product, and business unit.</li> <li> Cost allocation and chargeback / showback aligned to ownership and accountability.</li> <li> Rightsizing and usage optimisation across compute, storage, and AI workloads.</li> <li> Policy-based spend controls embedded into cloud operations.</li> </ul>
<b>Investment profile</b>	<ul style="list-style-type: none"> <li>• <b>Budget:</b> ~10–15% of total AI / tech spend.</li> <li>• <b>Timeline:</b> 3–6 months to establish visibility and begin measurable savings.</li> </ul>



## Success metrics



### Cost control & visibility

- 90–100% spend allocated.
- App-level cost visibility.
- Forecast accuracy  $\pm 10\text{--}15\%$ .



### Optimisation outcomes

- 20–40% cloud cost savings.
- 15–30% resource rightsizing.
- Reduced idle capacity.



### Operating discipline

- Named cost owners.
- Monthly FinOps reviews.
- Spend thresholds enforced.



### AI & scale readiness

- AI workloads cost-tracked.
- Unit economics defined.

## Market signals



### Australia Post

Uses Apptio Cloudability for cloud cost optimisation and FinOps-style cloud financial [management](#).



### Australian Government (AGA)

Recognises Cloud FinOps as a core capability for optimising public sector [cloud investments](#).

## Value creation use cases

### Finance

- Product chargeback
- Unit economics
- Spend controls



### Healthcare

- Cost transparency
- Compute optimisation
- Budget predictability



### Government / Public Sector

- Spend accountability
- Budget forecasting
- Shared services



### Utilities / Energy

- Cost control
- Workload forecasting
- Operational reliability



## Sample vendors

### Global

- **AWS:** cost explorer / savings plans
- **Microsoft:** Azure cost management
- **Google Cloud:** cloud billing / FinOps Hub

### Local

- **NCS Australia:** FinOps advisory services
- **Telstra Purple:** cloud cost optimisation services
- **Deloitte Australia:** FinOps services



The FinOps optimisation trap: Once costs are optimised, the question becomes where to invest next; using cloud savings to enable outcomes like data centre exit, scalability, and growth."

### Sidd Jolapara

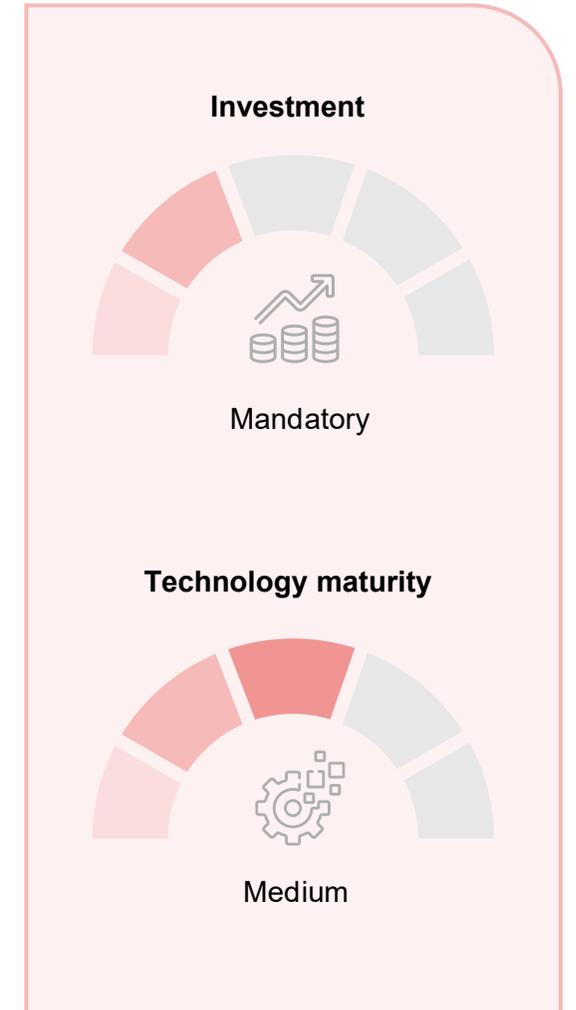
AWS' Head of Strategic Transformation for Public sector

[Read more](#) →

## Trend 4

# AI security guardrails

<b>Role</b>	<b>Protecting the enterprise while scaling AI</b>
<b>Definition</b>	A set of policies, controls, and monitoring mechanisms that ensure AI systems are used securely, responsibly, and in compliance with enterprise risk, privacy, and regulatory requirements.
<b>Why this matters now</b>	As GenAI adoption accelerates, security, data leakage, and misuse risks scale faster than traditional controls. Guardrails allow CIOs to enable AI usage without exposing the enterprise to unacceptable risk.
<b>Capabilities</b>	<ul style="list-style-type: none"> <li> Access controls and identity enforcement for AI tools and models.</li> <li> Data protection and prompt filtering to prevent sensitive data exposure.</li> <li> Usage monitoring and audit logging across AI interactions.</li> <li> Policy enforcement aligned to security, privacy, and compliance standards.</li> </ul>
<b>Investment profile</b>	<ul style="list-style-type: none"> <li>• <b>Budget:</b> ~8-12% of total AI / tech spend.</li> <li>• <b>Timeline:</b> 3–6 months to initial control; 6–12 months to enterprise coverage.</li> </ul>



## Success metrics



### Visibility & coverage

- **90–100%** models monitored.
- **80–95%** deployments automated.
- **60–80%** rollback within SLA.



### Control & enforcement

- **90–100%** access policy-controlled.
- **80–95%** prompts / data filtered.
- **75–90%** role-based enforcement.



### Risk reduction

- **70–90%** misuse detected.
- **60–80%** faster containment.



### Assurance & compliance

- **90–100%** audit log coverage.
- **80–95%** policy traceability.
- **100%** regulator-scoped systems covered.

## Market signals



Australian Government

### National AI Plan

Elevates governance and oversight expectations for organisations [adopting AI](#).



### NSW flood data breach

Contractor uploaded sensitive data to ChatGPT, exposing 3,000+ [individuals](#).

## Sample vendors

### Global

- **Microsoft:** Purview · Defender for cloud · Azure AI safety
- **AWS:** bedrock guardrails · IAM · CloudTrail
- **Zscaler:** AI data protection

### Local

- **Tesserent:** AI & data security services
- **Telstra Purple:** AI security & governance services
- **Vault Cloud:** sovereign secure AI environments

“

We're at the start of a golden era. Pre-trained models and agentic AI reduce the data barriers and allow cyber defence to operate more autonomously. From a technical perspective, I am extremely optimistic.”

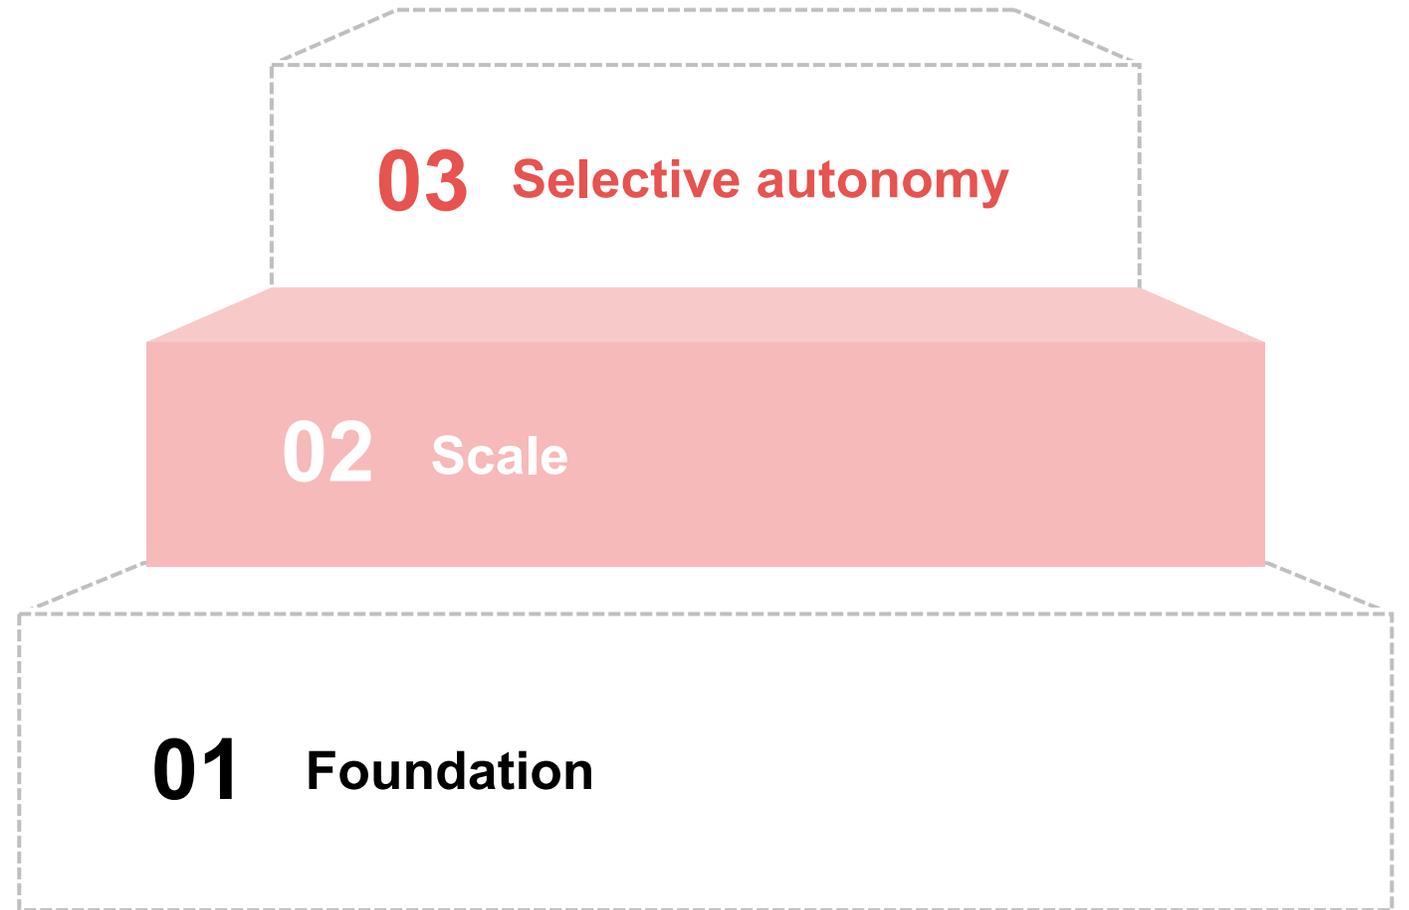
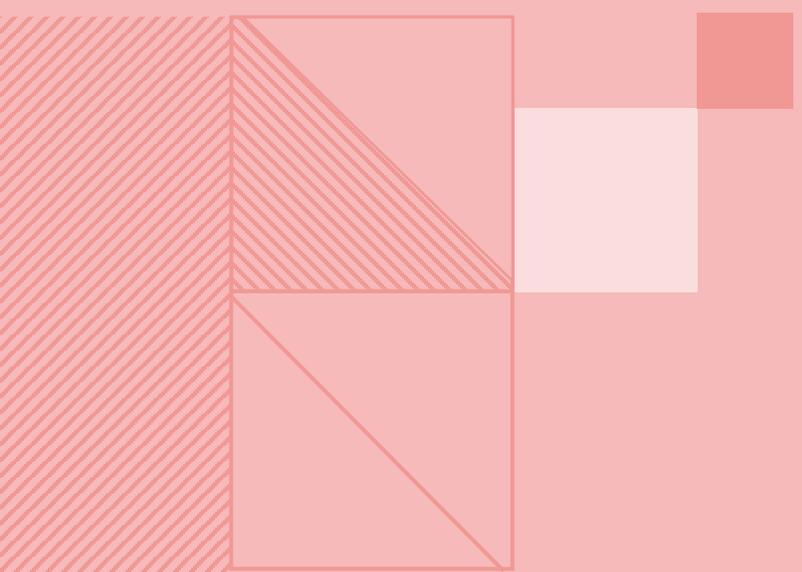
### William MacMillan

Former CIA CISO and now Chief Product Officer at Andesite

Read more [————>](#)

# 02

## Scale



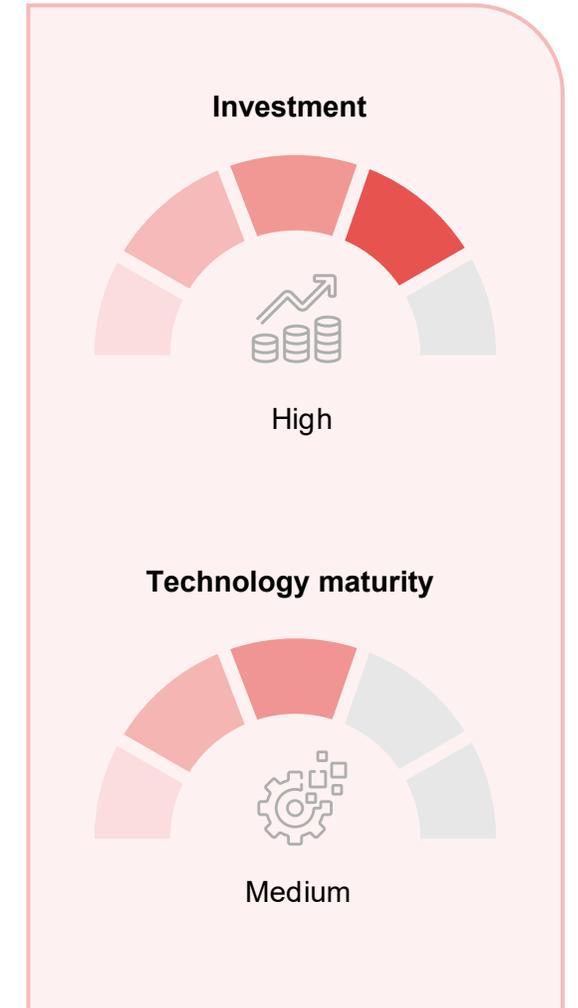
# 02

## Scale

### Trend 5

# MLOps & AI evaluation platforms

<b>Role</b>	<b>Scaling AI reliably and responsibly</b>
<b>Definition</b>	An enterprise capability that manages the deployment, monitoring, and evaluation of AI models to ensure performance, reliability, and compliance as AI moves from pilots into production.
<b>Why this matters now</b>	As organisations move beyond experimentation, AI risk shifts from “can it work?” to “can we trust and operate it at scale?” Without MLOps and evaluation, models degrade, risks go undetected, and AI initiatives fail to scale.
<b>Capabilities</b>	<ul style="list-style-type: none"><li> Model lifecycle management from development to production.</li><li> Continuous performance monitoring for drift, bias, and accuracy.</li><li> Evaluation and testing frameworks for GenAI and predictive models.</li><li> Versioning, auditability, and rollback controls for regulated environments.</li></ul>
<b>Investment profile</b>	<ul style="list-style-type: none"><li>• <b>Budget:</b> ~12-15% of total AI / tech spend.</li><li>• <b>Timeline:</b> 6–12 months to production-grade capability.</li></ul>



## Success metrics



### Operational stability

- **90–100%** models monitored.
- **80–95%** deployments automated.
- **60–80%** rollback within SLA.



### Model performance

- **70–90%** drift detected early.
- **85–95%** accuracy maintained.
- **60–80%** bias flagged pre-impact.



### Risk & assurance

- **90–100%** model traceability.
- **80–95%** audit evidence coverage.
- **100%** regulated models governed.



### Scale readiness

- **70–90%** pipelines reusable.
- **60–80%** teams standardised.

## Market signals



### AI misidentification error in Australian media

AI systems used in radio content led to a reporter being wrongly identified as a violent offender, triggering on-air corrections and internal [review](#).

## Sample vendors

### Global

- **Databricks:** MLflow
- **AWS:** SageMaker MLOps
- **Microsoft:** Azure machine learning
- **Google Cloud:** Vertex AI MLOps
- **IBM:** Watson OpenScale

### Local

- **DataRobot ANZ:** MLOps platform
- **Quantium:** model governance & evaluation
- **CSIRO's Data61:** trusted AI toolkits

## Value creation use cases



### Healthcare

- Model validation
- Bias monitoring
- Clinical assurance



### Utilities / Energy

- Forecast stability
- Anomaly detection
- Operational reliability



### Government / Public Sector

- Model transparency
- Performance oversight
- Policy compliance

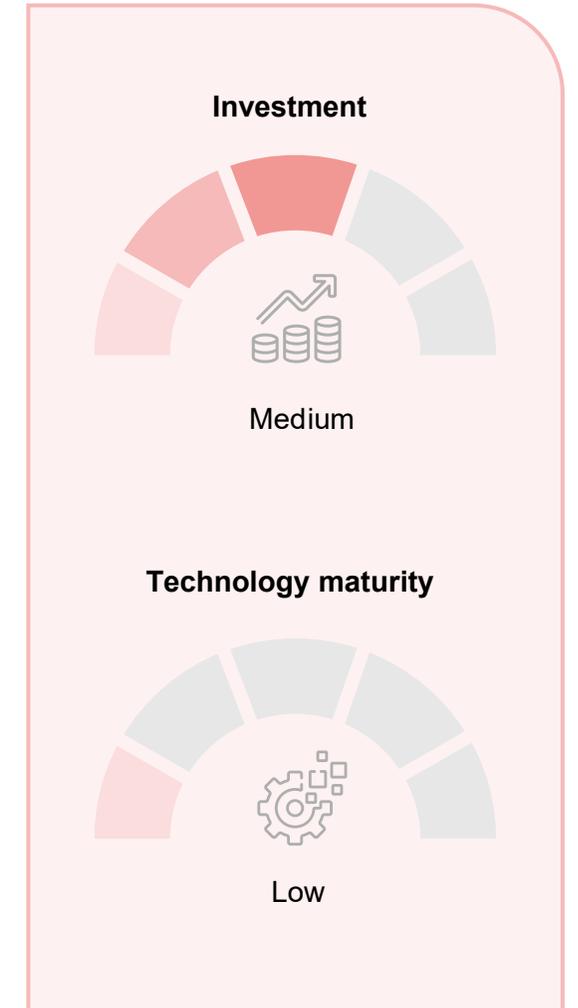
# 02

## Scale

### Trend 6

# Agentic AI orchestration platforms

<b>Role</b>	<b>Coordinating AI actions across enterprise workflows</b>
<b>Definition</b>	A control layer that coordinates multiple AI agents to plan, execute, and monitor multi-step tasks across systems, while keeping humans, policies, and controls in the loop.
<b>Why this matters now</b>	Enterprises are moving from AI insights to AI action. Without orchestration, agent-based AI becomes brittle, unsafe, and unscalable. Orchestration platforms provide the governance, sequencing, and control needed to operationalise agentic AI responsibly.
<b>Capabilities</b>	<ul style="list-style-type: none"> <li> Task orchestration across applications and data sources.</li> <li> Human-in-the-loop controls for approval and exception handling.</li> <li> Policy-driven execution aligned to security and compliance.</li> <li> Monitoring and logging of agent actions and outcomes.</li> </ul>
<b>Investment profile</b>	<ul style="list-style-type: none"> <li>• <b>Budget:</b> ~8-10% of total AI / tech spend.</li> <li>• <b>Timeline:</b> 6–12 months for controlled use cases; 12–24 months for broader autonomy.</li> </ul>



## Success metrics



### Operational control

- **90–100%** actions logged.
- **80–95%** tasks orchestrated.
- **<5%** uncontrolled failures.



### Business impact

- **20–40%** cycle-time reduction.
- **30–50%** manual effort reduction.
- **10–25%** throughput uplift.



### Risk & governance

- **90–100%** policy enforcement.
- **70–90%** human-in-loop coverage.
- **95–100%** audit visibility.



### Scale readiness

- **70–90%** pipelines reusable.
- **60–80%** teams standardised.

## Market signals



**Salesforce** launched Agentforce 360, enabling enterprises and partners to build, orchestrate, and govern AI agents across CRM and business [workflows](#).



**UiPath** and **Talkdesk** are combining agentic AI to automate healthcare contact-centre workflows, improving efficiency and patient [experience](#).

## Value creation use cases

### Finance

- Client onboarding automation.
- Regulatory response flows.
- Multi-system reconciliation.



### Healthcare

- Claims processing automation.
- Patient service orchestration.
- Documentation workflows.



### Retail

- Customer experience bots.
- Order fulfillment orchestration.
- Campaign automation.



## Sample vendors

### Global

- **Salesforce:** Agentforce 360
- **Oracle:** AI Agent Studio
- **UiPath:** AI orchestration / AI centre
- **Microsoft:** Copilot Studio / Autogen

### Local

- **Deloitte Australia:** agentic AI & intelligent automation services
- **Servian:** AI workflow orchestration services
- **Mantel Group:** AI workflow & automation architecture



Reusable platforms are the strongest predictor of profitability and innovation, shifting value away from projects and into enduring business assets.”

### Bruce Northcote

Chief Compliance and Chief Security Officer at University of Adelaide

Read more [→](#)

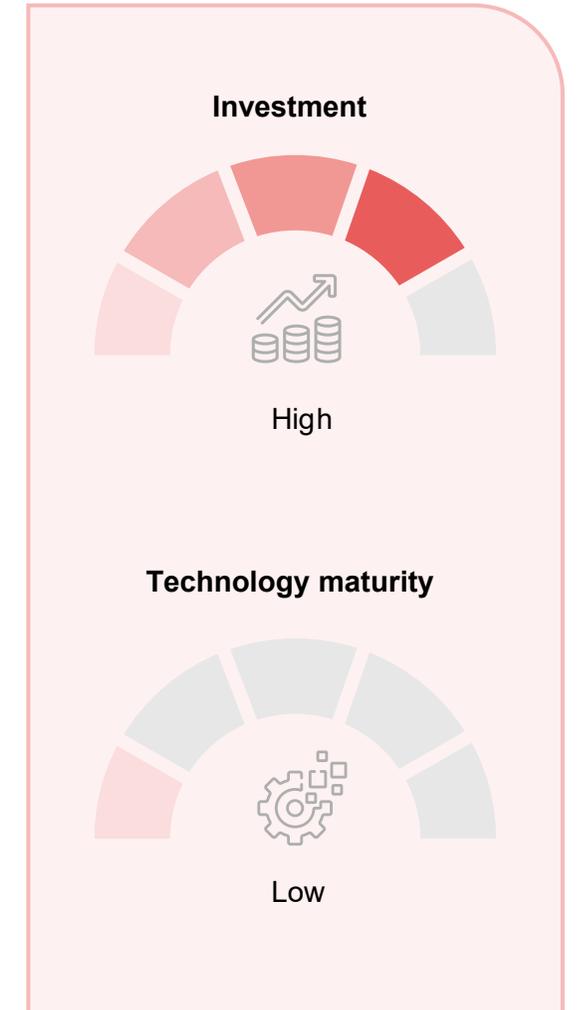
# 02

## Scale

### Trend 7

# ERP modernisation with AI support

<b>Role</b>	<b>Making core systems more intelligent, efficient, and adaptable</b>
<b>Definition</b>	The modernisation of core ERP platforms (SAP, Oracle, Workday) using AI to improve process efficiency, decision quality, and user experience, without destabilising mission-critical systems.
<b>Why this matters now</b>	Legacy ERP environments are costly, rigid, and slow to adapt. AI enables organisations to extract more value from ERP investments while reducing manual effort, improving forecasting, and supporting ongoing cloud and operating-model transformation.
<b>Capabilities</b>	<ul style="list-style-type: none"><li> Process automation across finance, supply chain, HR, and procurement.</li><li> AI-assisted forecasting and planning using ERP data.</li><li> Intelligent user support embedded into ERP workflows.</li><li> Faster upgrades and reduced technical debt during ERP transitions.</li></ul>
<b>Investment profile</b>	<ul style="list-style-type: none"><li>• <b>Budget:</b> ~15-20% of total AI / tech spend.</li><li>• <b>Timeline:</b> 6–12 months for priority processes; 12–24 months enterprise-wide.</li></ul>



## Success metrics



### Process efficiency

- 20–40% cycle-time reduction.
- 15–30% manual effort reduction.
- 10–25% exception reduction.



### Decision & planning quality

- 10–20% forecast error reduction.
- 20–40% faster planning cycles.
- 15–30% fewer overrides.



### Control & stability

- 99.5–99.9% process uptime.
- 90–100% automation governed.



### Modernisation momentum

- 20–40% faster upgrades.
- 15–30% less custom code.
- 60–80% processes standardised.

## Market signals



**Annexa** adds AI innovations to NetSuite ERP across Australia and [New Zealand](#).



**Infosys** collaborates with Melbourne Archdiocese Catholic Schools to drive Network-Wide Digital [Transformation](#).



**Kyndryl & Nova Intelligence** accelerate SAP Cloud ERP modernisation with [AI tooling](#).

## Sample vendors

### Global

- **ServiceNow:** Now Assist, AI workflows
- **Salesforce:** Einstein, MuleSoft for ERP integration
- **IBM:** Watsonx, consulting ERP AI accelerators

### Local

- **Accenture Australia:** Intelligent ERP transformation
- **Deloitte Australia:** AI-enabled ERP modernisation
- **PwC Australia:** ERP cloud & AI enablement
- **Kyndryl Australia:** SAP cloud & AI modernisation



Reusable platforms are the strongest predictor of profitability and innovation, shifting value away from projects and into enduring business assets.”

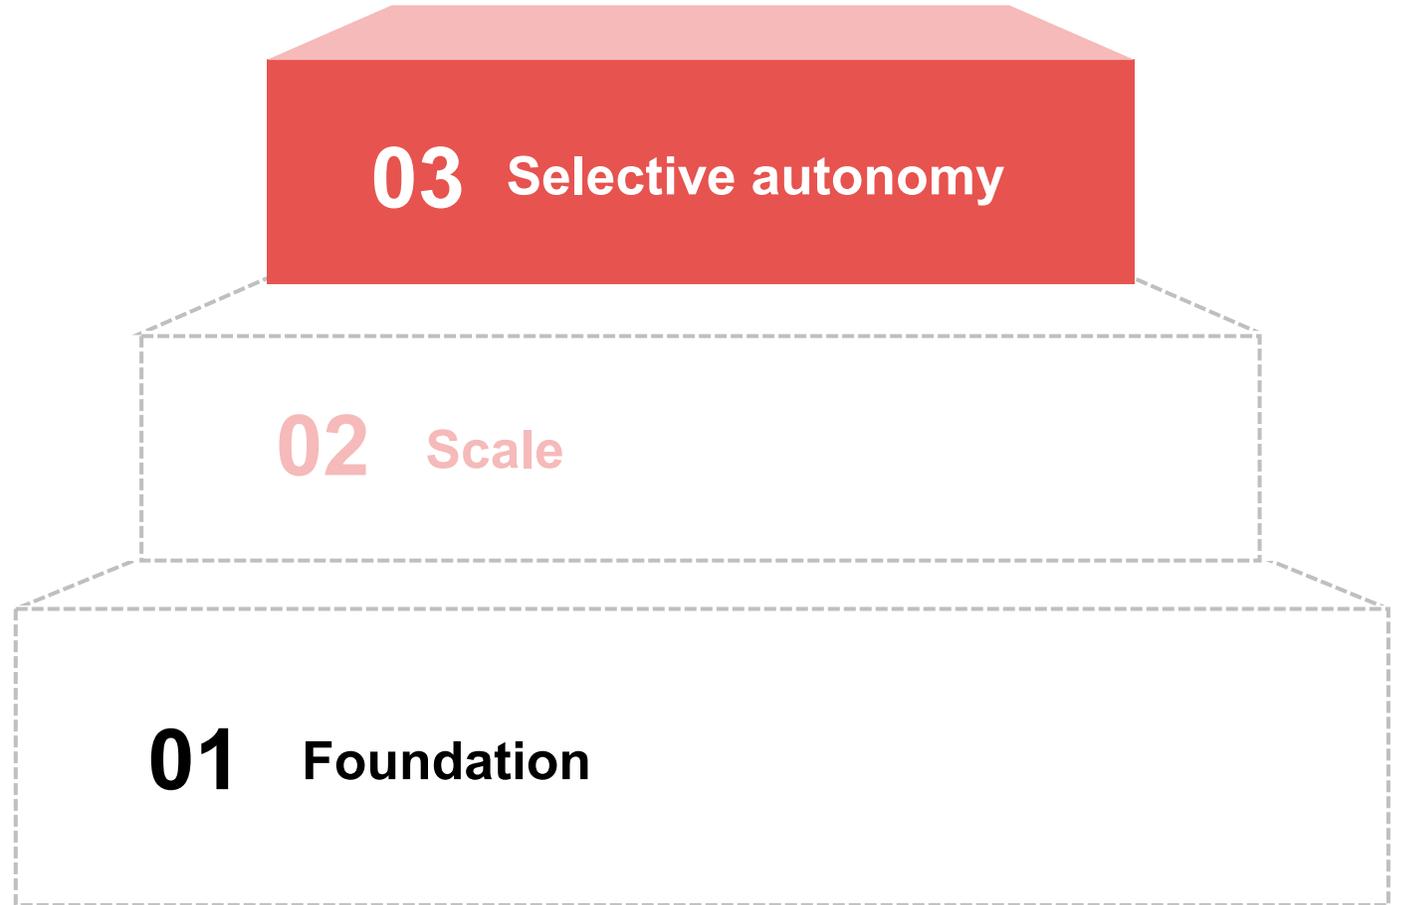
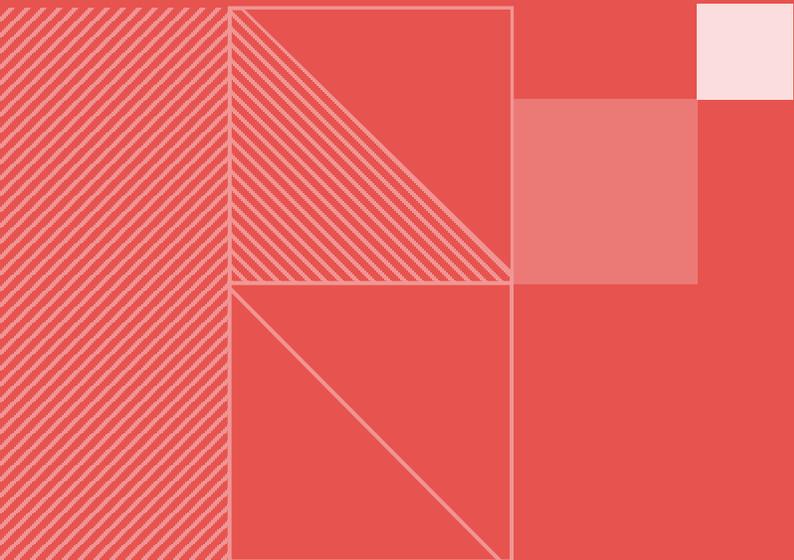
### Peter Weill

CISR Chairman & Senior Advisor at ADAPT

Read more 

03

# Selective autonomy



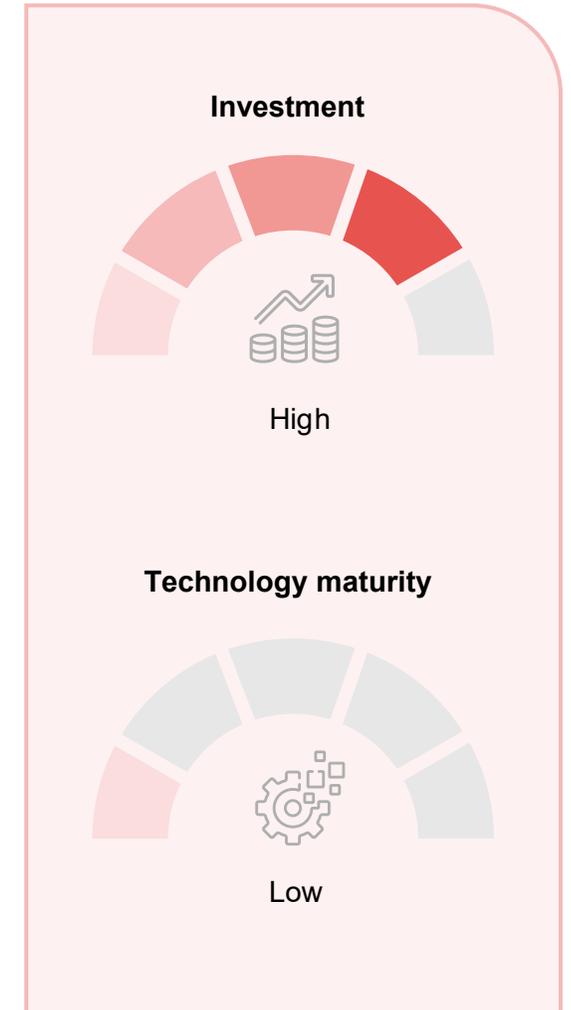
# 03

## Selective autonomy

### Trend 8

# Real-time data fabric (event mesh)

<b>Role</b>	Enabling high-density compute for AI, sustainability, and cost control
<b>Definition</b>	A technology-driven architecture that unifies data access across disparate systems (Data Fabric) and integrates a dynamic infrastructure layer for distributing events (Event Mesh) to enable immediate, event-driven processing.
<b>Why this matters now</b>	Traditional batch-oriented data architectures and centralised data lakes create bottlenecks, hindering agility and the ability to act on data-in-motion. The demand for AI/ML models and personalised customer experiences requires instant, governed access to all enterprise data.
<b>Capabilities</b>	<ul style="list-style-type: none"> <li> Real-time data availability (near-zero latency access to data-in-motion).</li> <li> Automated data discovery (metadata-driven integration and knowledge graphs).</li> <li> Decoupled event routing (scalable, reliable event distribution via Event Mesh).</li> <li> Unified data governance (consistent security and quality policies across all sources).</li> </ul>
<b>Investment profile</b>	<ul style="list-style-type: none"> <li>• <b>Budget:</b> ~5-8% of total AI / tech spend.</li> <li>• <b>Timeline:</b> 6–12 months for priority processes; 12–24 months enterprise-wide.</li> </ul>



## Success metrics



### Operational efficiency

- Reduced time-to-insight (e.g., 50% faster).
- 30% reduction in data integration design / deployment time.
- Lower data management costs.



### Business agility

- Faster time-to-market for new data products.
- Increased number of real-time applications.
- Improved decision-making speed.



### Data governance

- Consistent enforcement of security and quality policies.
- Improved data lineage and compliance scores.
- Reduced data quality incidents.



### Performance

- Reduced data latency (near-zero).
- Increased event throughput (events / second).
- High system uptime and reliability.

## Market signals



### Australian Red Cross

**Australian Red Cross** adopts data mesh strategy powered by [Fabric](#).



### Banking case study

Modernisation of a monolithic payment processing system to a microservices-based, event-driven architecture using Solace (an Event Mesh [technology](#)).

## Sample vendors

### Global

- **Microsoft:** Microsoft Fabric
- **Confluent Platform**
- **Solace**

### Local

- **Datafabric:** Compass
- **Zetaris:** Modern Lakehouse for AI
- **Intersystems:** Data Studio



What we're hoping to get out of Fabric is an understanding of how we can bring together different data sets from different sources and view the data in ways we haven't been able to before."

### Brett Wilson

CIO at Australian Red Cross

Read more [————>](#)

# 03

## Selective autonomy

### Trend 9

# Data centre evolution to liquid cooling

<b>Role</b>	Enabling high-density compute for AI, sustainability, and cost control.
<b>Definition</b>	A shift from traditional air cooling to direct-to-chip, immersion, or hybrid liquid cooling to support high-density AI workloads, improve energy efficiency, and extend data-centre capacity without major footprint expansion.
<b>Why this matters now</b>	AI, HPC, and accelerated workloads are exceeding the limits of air-cooled infrastructure. Liquid cooling is now required to unlock GPU density, manage power costs, and meet sustainability targets without building entirely new data centres.
<b>Capabilities</b>	<ul style="list-style-type: none"><li> Higher compute density per rack.</li><li> Lower energy consumption and cooling overhead.</li><li> Improved performance stability for AI workloads.</li><li> Extended lifespan of existing data-centre assets.</li></ul>
<b>Investment profile</b>	<ul style="list-style-type: none"><li>• <b>Budget:</b> ~15-20% of total AI / tech spend.</li><li>• <b>Timeline:</b> 24 months for priority processes; 12–24 months enterprise-wide.</li></ul>



## Success metrics



### Energy & cost efficiency

- 20–40% cooling energy reduction.
- Lower PUEP.
- Power cost stabilisation.



### Compute enablement

- 2–5x rack density.
- Stable GPU performance.
- Reduced thermal throttling.



### Sustainability

- Lower carbon intensity.
- Heat reuse enabled.
- ESG targets supported.



### Infrastructure readiness

- AI-ready facilities.
- Incremental retrofits.
- Scalable cooling design.

## Market signals



N E X T D C

**NEXTDC** announces major AI Factory investment (AUD 2 billion) to support high-density AI and cooling [innovation](#).



**DXN** launches liquid-cooled AI / HPC data centre [module](#).

ResetData



**ResetData + BP Castrol** partner on AI-optimised with liquid [cooling](#).

## Sample vendors

### Global

- **Dell Technologies:** PowerEdge liquid-cooled servers
- **NVIDIA:** DGX liquid-ready platforms
- **HPE:** Cray / GreenLake for HPC
- **Supermicro:** liquid-cooled AI systems

### Local

- **Canberra Data Centres:** HPC & sovereign facilities
- **Macquarie Data Centres:** high-power colocation
- **DXN Limited:** liquid-cooled modular data centres
- **NEXTDC:** high-density AI-ready facilities

“

The rapid growth of AI is driving unprecedented power and cooling demands that existing data centre designs can no longer support.”

**Matt Gurr**

ANZ Market Director – High Tech, AECOM

[Read more](#) →

“

Over the next three years, data centres are being designed for up to 1,000 kilowatts per rack, fully liquid-cooled and purpose-built for AI.”

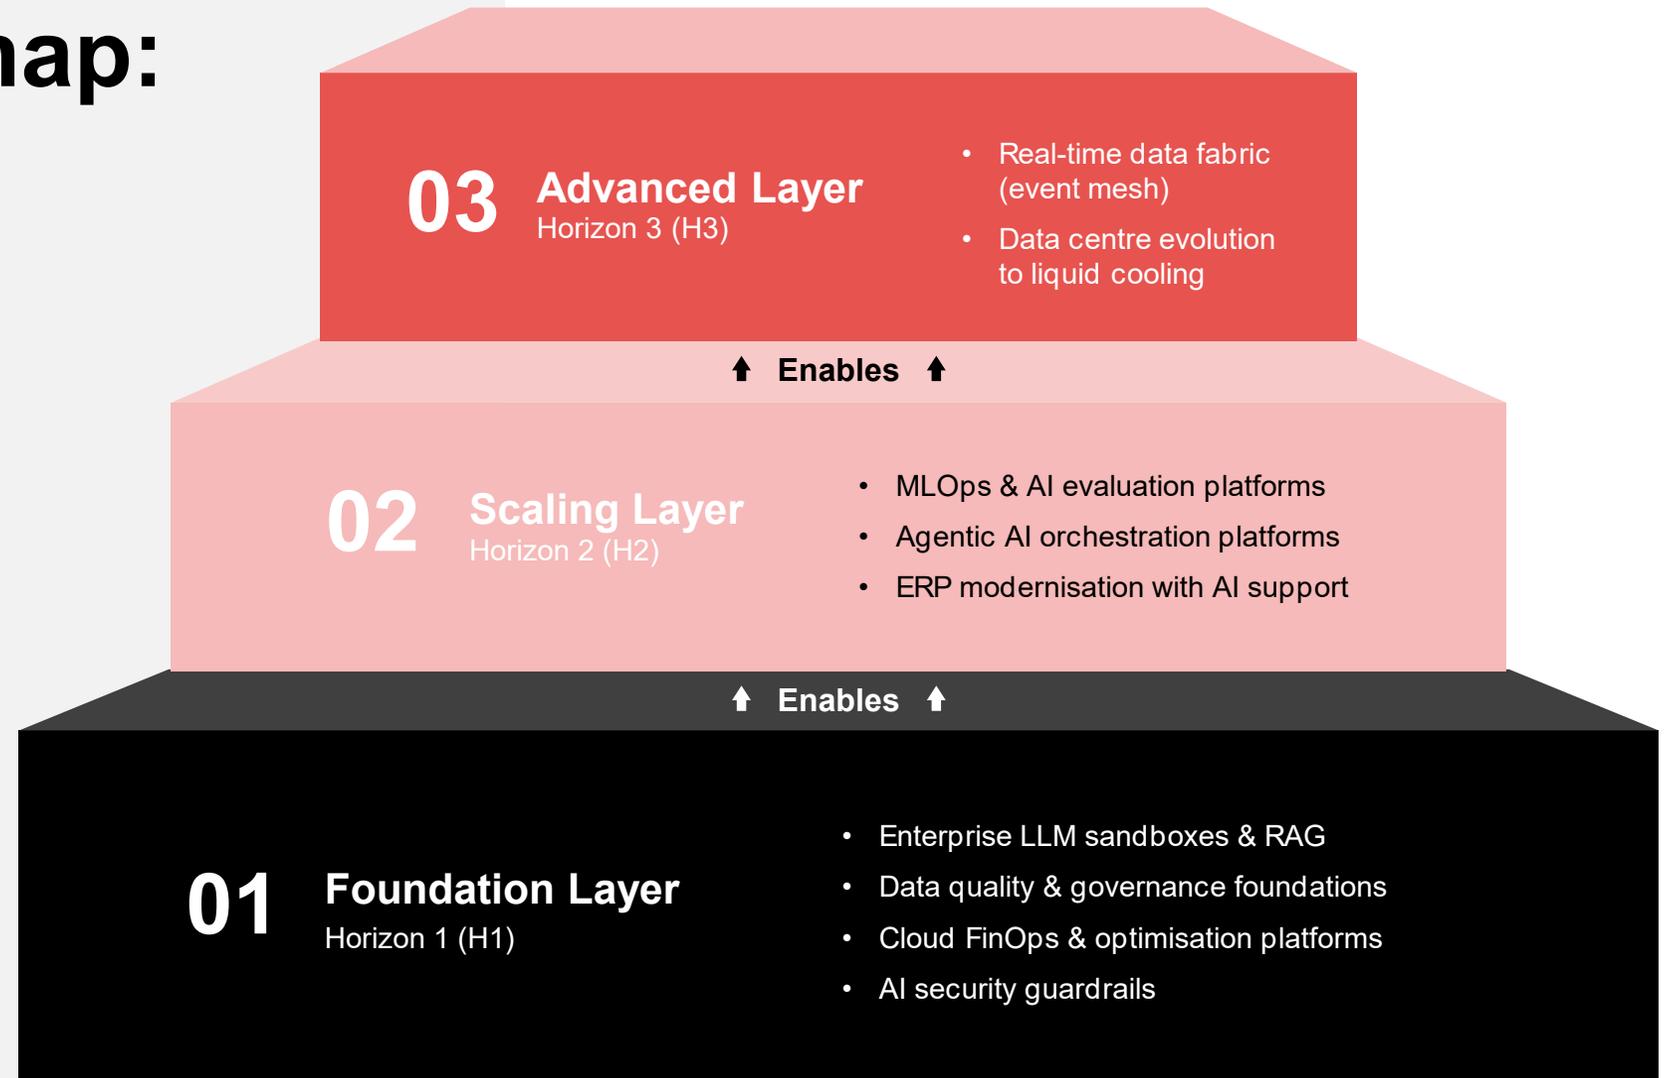
**Raj Singh**

Targeted Segments Enterprise Sales Leader at Schneider Electric

[Read more](#) →

# Technology dependency map:

## You cannot skip layers



# Quick reality check

## Questions to validate your journey

### Ask today

- AI foundations:** How much of our AI budget goes to Horizon 1 vs scaling?
- Data readiness:** Do we have owners and governance for our top 10 AI use cases?
- Security and risk:** Are all GenAI interactions logged, monitored, and policy-controlled?
- Cost control:** Can we link cloud spend to business outcomes and unit costs?

### Ask in the next 30 days

- Investment sequencing:** Are Horizon 1 dependencies confirmed before we fund Horizon 2 or 3?
- Gen AI enablement:** What is our 90-day adoption target for the enterprise LLM sandbox?
- FinOps:** Who owns cloud cost accountability per application, and when do monthly reviews start?
- AI governance:** Do 90%+ of our AI models have monitoring, access controls, and audit logs?
- MLOps readiness:** Can we detect model drift, bias, and performance issues in production?

# How ADAPT supports you



# Leveraging ADAPT in your transformation journey



## Validate your priorities

- Map current initiatives to the technology shifts in this report.
- Check for gaps between current investment and business-critical use cases.
- Review how spending is split between foundational and emerging capabilities.



## Pressure-test your roadmap

- Review sequencing against maturity, risk, and time-to-value.
- Assess success metrics against peer benchmarks and internal targets.
- Identify work that needs tighter scope or re-phasing.



## Engage with ADAPT advisors

- Test assumptions with independent CIO-peer insights.
- Turn findings into a prioritised, executable roadmap.
- Ensure leadership and delivery teams are aligned on clear outcomes.

# Advisors-on-Demand



**Delan Naidu**  
Former Tech Strategy  
Lead, Qantas & TransGrid



**Dr. Patrick Maes**  
Former Group CTO, ANZ  
& Global CIO, Credit Suisse



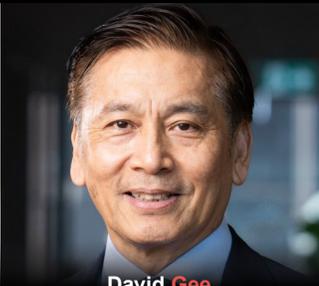
**David Heacock**  
Former Chief Digital Innovation, Officer  
at DXC Technology & BCG Partner



**Sally Osborne**  
Former National Director Colliers  
and COO Belle Property Group



**Sami Yalavac**  
Former CIO & MD at Bupa



**David Gee**  
CIO, CISO,  
Risk Executive & Author



**Ellis Brover**  
Former CIO of Toyota



**Mark Pesce**  
Co-Founder at Wisely AI



**Jennifer Scott**  
Former COO Ebay & GM Digital  
Transformation ANZ Bank



**Nathan Moore**  
Former Chief Nursing  
Information Officer, NSW Health



**Marcus Darbyshire**  
Former CIO, South East Water &  
Gartner Executive Partner



**Archer Reed**  
Former CTO, Hewlett Packard



**Sheridan Ware**  
Former CITO at Charter Hall



**Simon Kriss**  
Chief AI Officer, simonkriss.ai



**Andy Foster**  
Principal & Executive Coach,  
Clarity Partners



**Claudine Ogilvie**  
Former CIO at Jetstar & APAC Tech  
Executive for Compass Group



**Peter Tsakissiris**  
Former Chief Digital Officer, DXC  
& Gartner Executive Partner



**Michael Smart**  
Former KPMG Partner & Gartner  
Consulting VP



**Sarah Hughes**  
Former, Chief Nursing  
Information Officer, Olinqua



**Kate Rezenbrick**  
Former Chief Nursing & Midwife  
Information Officer



**Samantha McIntyre**  
Founder Techboosters. Former CIO  
Target, Sussan Group & L'oreal



**James Foxall**  
Former CIO, Metro South Health



**Memo Hayek**  
Former CIO at  
NAB and Commbank



**Rowan Dollar**  
Former CIO, Catholic Education,  
Archdiocese of Canberra & Goulburn



**Nitesh Chovhan**  
Former Head of eBusiness Tech  
& Strategy, Nestle



**Mark Cameron**  
CEO of Alyve



**Rosalie Adriano**  
Former CIO of Nestlé Oceania



**Alan Sharvin**  
Former CIO of Tabcorp



- Governance, risk & compliance
- Cyber security & resilience
- Technology strategy & execution
- Data Governance
- Executive & stakeholder management

# Claudine Ogilvie

**Former CIO at Jetstar & APAC Tech Executive for Compass Group**

<https://www.linkedin.com/in/claudine-ogilvie/>

Claudine is an accomplished entrepreneur, CEO, and board director with 20 years of leadership across strategy, innovation, and emerging technologies. She brings deep expertise in data, AI, quantum computing, and cyber governance, with a track record spanning aviation, finance, insurance, and consulting. Known for blending commercial acumen with visionary thinking, Claudine helps organisations navigate complex technology landscapes with confidence.

- **Managing Director at O&O Consulting and CEO of HivePix.**
- **Board Director at Cuscal, Scyne Advisory, and formerly Youi Insurance.**
- **Former CIO at Jetstar Group and Lead technology, data businesses and innovation at Compass Group APAC.**
- **Past leader of Consumer and Industrial Markets at KPMG.**
- **Recognised in Australia's Top CIO50 (2016).**
- **Holds qualifications from MIT (Quantum Computing), GAICD, UTS, and ESC Reims.**
- **Active adviser and coach to leaders in innovation, diversity, and ESG.**



- AI & data management
- Technology strategy & execution
- Executive & stakeholder management
- Governance, risk & compliance
- Leadership development

# Mark Cameron

**CEO & Director, Alyve**

<https://www.linkedin.com/in/markrcameron/>

Mark is the founder and CEO of Alyve, a leading digital transformation consultancy. He is known for integrating AI and emerging tech into enterprise strategy, with a focus on improving customer experience and operational outcomes. Mark brings deep expertise in digital strategy, innovation, and executive advisory.

- **CEO of award-winning digital transformation firm Alyve.**
- **Certified Experience Economy Expert.**
- **Specialist in AI-driven business and customer transformation.**
- **Works with senior executives on sustainable growth strategies.**
- **Frequent speaker and thought leader on digital innovation.**
- **Expertise in ESG-aligned technology and digital ethics.**



- Digital and operating model transformation
- Innovation and growth
- Human centred design
- Leadership development and coaching

# Sheridan Ware

**Former CITO, Charter Hall**

<https://www.linkedin.com/in/sheridanware/>

Sheridan is a multi-award-winning transformation leader with over 20 years' experience driving strategic digital initiatives focused on real estate and financial services. She partners with boards and executives to align emerging technology with purpose, creating value for customers, employees, and stakeholders through human-centred and sustainable digital design.

- **Former Chief Information & Technology Officer at Charter Hall.**
- **Previously CIO APAC at Cushman & Wakefield.**
- **Chair of the Sustainable Digitalisation Project.**
- **Specialises in human-centred transformation and digital engagement.**
- **Recognised for leadership in innovation, inclusion, and ESG-aligned technology.**
- **2023 recipient of Australia's CIO of the Year (CIO50 Awards).**



- AI & data management
- Technology strategy & execution
- Change management
- Business strategy & implementation
- Talent strategy

# Mark Pesce

**Co-Founder, Wisely AI**

<https://www.linkedin.com/in/markpesce/>

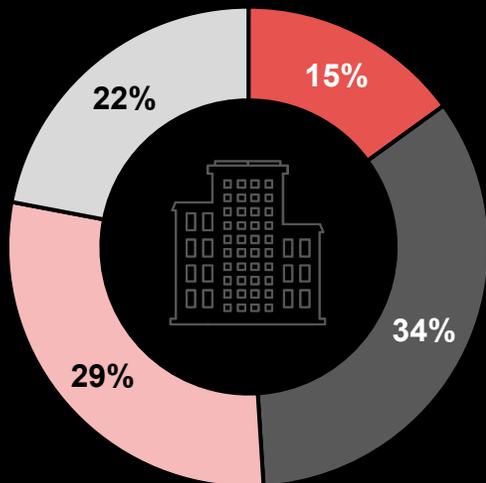
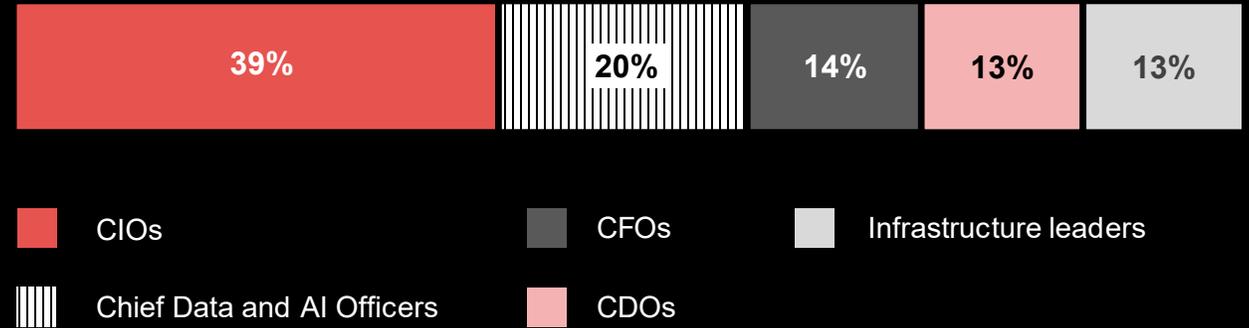
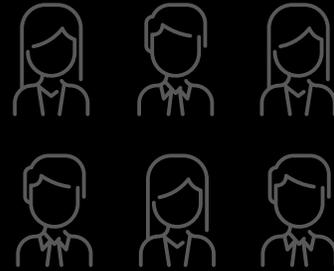
Mark is a futurist, inventor, and thought leader in emerging technology. Best known as the co-inventor of VRML (early web 3D tech), he has decades of experience helping organisations anticipate and navigate disruptive change - bringing a unique lens to AI, the future of work, digital finance, education, and media.

- **Author of nine books, including on AI and the metaverse.**
- **Adviser to organisations including the G20, World Bank, and Telstra.**
- **Host of award-winning tech podcasts and keynote speaker.**
- **Expert in AI adoption, digital finance, energy, and lifelong learning.**
- **Works with C-suite and educators on the future of work and digital literacy.**
- **A background in building capacity for disruption across sectors.**

# Survey demographics



974  
respondents



- **Small and medium organisations**  
Up to 500 employees
- **Mid-sized organisations**  
501 to 2,500 employees
- **Enterprise-sized organisations**  
2,501 to 10,000 employees
- **Large enterprise-sized organisations**  
More than 10,000 employees



**44%** of Australia's  
**A\$1.9t** GDP

Combined unique organisational revenues of over **A\$821b**  
See methodology for further details



**13%** of the Australian  
workforce

Combined unique workforce employed of over **1.7m**

## Written by



**Pooja Singh**  
**Senior Research Analyst at ADAPT**

Pooja is a subject matter expert in AI, data analytics, and emerging technologies, advising C-suite and product leaders across AI adoption strategy, AI use-case prioritization, and business-case development. She brings nearly a decade of experience across U.S. and Australian markets, shaped by her roles at Gartner (Principal Analyst) and Bain & Co. Pooja holds an engineering degree in biotechnology and an MBA from Johns Hopkins University, with a focus on translating AI innovation into measurable business value.



**Byron Connolly**  
**Head of Programs and Value Engagement at ADAPT**

Byron is a highly experienced technology and business journalist, editor, corporate writer, and event producer. Prior to joining ADAPT, he was the editor-in-chief at CIO Australia and associate editor at CSO Australia. He also created and led the well-known CIO50 awards program in Australia and The CIO Show podcast. Byron creates valuable insights for our community of senior technology and business professionals that help them reach their organisational and professional goals.



**Harshit Goel**  
**Senior Analyst, Strategic Research at ADAPT**

Harshit's role as a Senior Analyst at ADAPT focuses on delivering actionable insights that help technology and business leaders navigate emerging trends, optimise technology investments, and drive transformative value across industries. Specialising in the BFSI, fintech, telecom, and retail sectors, Harshit advises C-suite executives and business leaders. His expertise spans areas in digital banking and payments, artificial intelligence applications, and digital commerce and experience platforms.

## About ADAPT

ADAPT exists to help A/NZ's leading organisations become more agile, innovative, and resilient in the face of change.

We support IT and business leaders with exclusive access to A/NZ's largest executive community, 60,000+ local data points, and practical advisory all built on regional realities, not global generalisations.

ADAPT equips executive teams with the insights, alignment, and credibility to lead enterprise-wide transformation. We help you benchmark progress, influence stakeholders, and make decisions that deliver measurable and lasting impact.

# ADAPT

[adapt.com.au](https://adapt.com.au) | [hello@adapt.com.au](mailto:hello@adapt.com.au)



### Data Validation

Results contained in this report (Report) are based on survey responses as provided by our clients and publicly available information. ADAPT does not subject the data contained in the Report to audit or review procedures or any other testing to validate the accuracy or reasonableness of the data provided by the participating clients. While we do not manipulate the survey responses, we do make an effort to identify outlier data or conflicting results that could lead to misinterpretations before the preparation of the Report.

### Intended Use of the Report

The Report is designed to help technology executives in Australia and New Zealand make independent decisions regarding financial, resourcing and operational performance matters. The information and data contained in the Report are provided as is and are for information purposes only. They are not intended or implied to be recommendations and in no event will ADAPT be liable for any decisions or actions taken in reliance of the results obtained through the use of the information and/or data contained in the Report.

### Research MSA and Privacy Policy

The Report and any other information provided by ADAPT is subject to the Research and Advisory Master Services Agreement and the ADAPT Privacy Policy.