The next leap in trusted digital connectivity

Beyond 2025

Produced by

InfoTrack

1.

Executive summary

By 2026, Australia is entering a new epoch in digital transformation:
The rise of autonomous, compliance-aware AI agents connected directly to verified government data.

Agentic AI, systems that can reason, act, and collaborate across enterprise and government boundaries, will underpin the next decade of national productivity yet, enterprise adoption lags behind potential. According to the National AI Readiness Index 2025, 92% of businesses use generative AI, but only 19% use agentic AI, accounting for just 1% of value delivered.

The challenge is not capability, it's trust, compliance and integration.

InfoTrack's MCP Platform architecture bridges this gap, providing:

- **Verified data** provenance from official government sources
- **Enterprise-grade** auditability and privacy controls
- I**nteroperability** across Salesforce, Microsoft, OpenAI and government registries
- A **unified** trust layer that lets AI agents act with context, confidence, and compliance

This whitepaper outlines the strategic, regulatory and technological imperatives shaping AI adoption through to beyond 2026, and presents a blueprint for leaders in technology leading Australia's transition into an agentic AI economy.

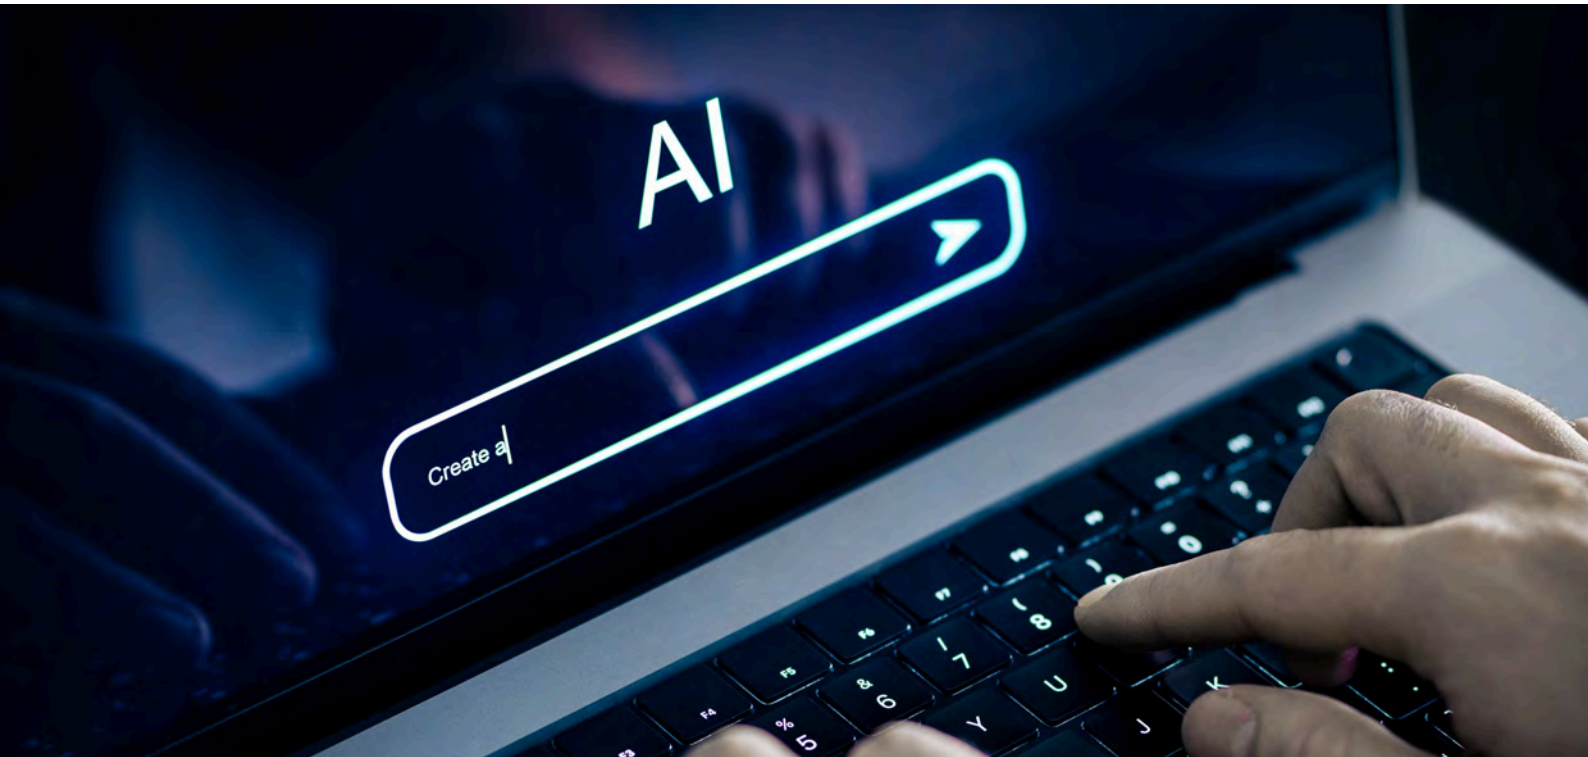

2.

From generative to agentic: A shift to autonomous intelligence

Generative AI has boosted individual productivity.
Agentic AI will transform **organisational and sovereign productivity.**

What is Agentic AI?

Agentic AI is compromised of autonomous systems capable of:

- Reasoning in line with organisational goals
- Executing end-to-end workflows
- Interacting securely with structured and unstructured data
- Making informed decisions based on verified information

This evolution mirrors the market's trajectory. According to Striim, 95% of enterprise AI pilots fail to reach production, not due to model quality but due to brittle integrations and lack of contextual data access.

Agentic AI solves this via:

- Real-time context
- Multi-system interoperability
- Autonomous orchestration
- Event-driven decisioning

"By automating complex business workflows, agents unlock the full potential of vertical use cases. Forward-looking companies are already harnessing the power of agents to transform core processes."
- McKinsey's 2025 Seizing the agentic AI advantage report

3.

Australia's AI landscape (2026 and beyond): Regulation, risk and opportunity

Australia's digital readiness has accelerated dramatically.

Key signals according to Department of Industry, Science and Resources, June 2025:

- According to the National Artificial Intelligence Centre (NAIC), 40% of Australian SMEs are currently adopting AI (as of Q4 2024), marking a 5% increase compared to the previous quarter.
- Among the smallest businesses (those with up to 4 employees), AI adoption rose from 25% to 34% over the same period.
- The proportion of SMEs reporting that they are "not aware of how to use AI" fell to 21%.
- Per the NAIC's latest report published in mid-2025, Australia now counts over 1,500 AI-focused companies contributing to the national AI ecosystem, illustrating growth in both adoption and domestic AI capability.

And yet one barrier consistently blocks progress:

"How do we adopt AI agents without breaking compliance, security or trust?"

The same Readiness Index reveals:

- **52%** of businesses say AI will be difficult to integrate
- Security & compliance are the #1 barrier to adoption (**29%)**

This "adoption paralysis" is exactly the gap that InfoTrack's MCP Platform solves, providing:
A trusted, compliant bridge between enterprise agents and government systems.

4.

The trust layer: InfoTrack's MCP Platform architecture

The Model Context Protocol (MCP) is becoming the USB-C of enterprise AI connectivity, a universal standard that allows any AI agent to securely connect to any data source.

InfoTrack extends MCP with Australia's most proven government-grade trust infrastructure.

What InfoTrack Provides:

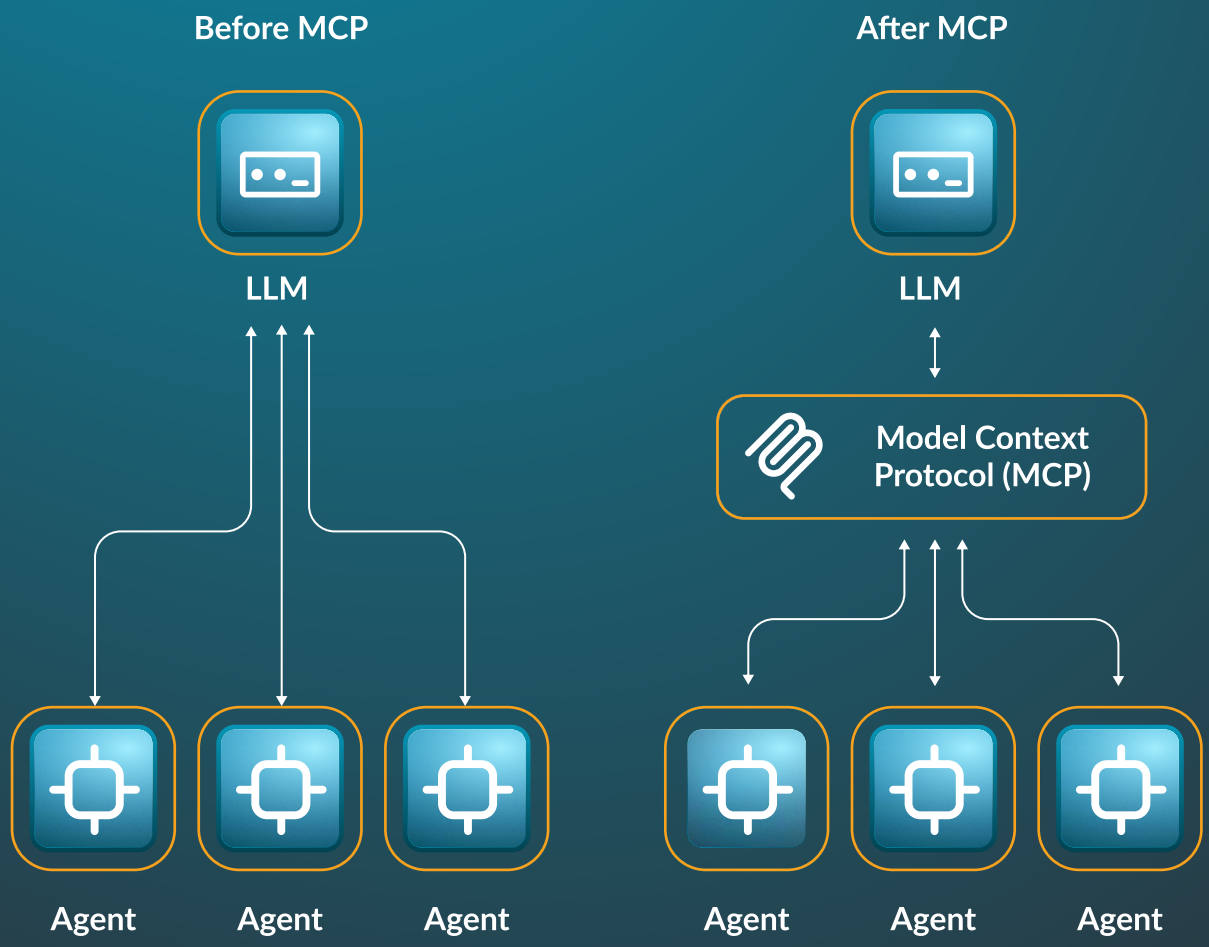

Enabling access and reassurance through government verification, built-in compliance and secure integrations.

A. Verified Government Data

Direct, certificate-based access to:

- ASIC
- Land Titles & Property Registries
- Courts
- PPSR

B. Compliance-by-Design

- Full audit trails of every lookup, action & decision
- ISO 27001-aligned security
- Permission-controlled access
- Privacy-safe data routing

C. Secure MCP Integration

MCP removes brittle, hard-coded integrations. Instead:

- AI agents request only what they're authorised to access
- A central MCP server mediates all interactions
- Government data never interacts directly with an AI model
- Sensitive sources remain protected behind verified trust frameworks

5.

Cross-industry use cases (2026–2030)

InfoTrack's MCP Platform Agentic AI yields the highest return across these sectors.

A. Legal & Conveyancing

2026–2028 Transformation

- Automated title searches, ASIC extracts & court filings
- AI-assisted matter intake, risk checks and conflict analysis
- Real-time settlement readiness verification

Impact:

- Pilot modelling shows 60% faster settlement cycles using AI-enabled certificate retrieval and validation.

B. Banking, Finance & Insurance

The Striim report highlights that banks need real-time 360 customer views and compliance-safe access to operational data.

Example scenarios

- Autonomous KYC/AML workflows
- AI agents that perform real-time title verification for secured lending
- Automated portfolio monitoring
- Embedded fraud detection powered by continuous MCP-enabled data streams

Expected outcomes (Deloitte 2025):

25–35% reduction in manual compliance workloads.

C. Government, Regulators & Public Safety

Enabled through Permissioned MCP Sandboxes

- AI-guided license applications
- Multi-agency event-driven compliance triggers
- Real-time digital identity and credential verification
- Autonomous report population and audit preparation

Government adoption of agentic systems accelerates when **data provenance is guaranteed.**

D. Health, Aged Care & Social Services

From Striim's use cases (page 12):

- AI agents monitoring patients and cross-checking with EHRs in real time
- Early anomaly detection through secure event streams
- Automated compliance reporting for Medicare & NDIS audits

Impact:
Reduces clinical workload while improving treatment outcomes.

E. Commercial & Enterprise Operations

- AI agents for HR compliance verification
- Procurement agents validating ABNs/ACNs on the fly
- Enterprise risk agents identifying anomalies across government data feeds
- Zero-touch onboarding of suppliers, contractors, and partners

6.

Risk and compliance architecture for agentic AI

Enterprises adopting AI agents face three systemic risks:

1. **Unverified data → Unreliable decisions**
2. **Opaque workflows → Regulatory exposure**
3. **Direct database access → System instability & security risk**

Striim notes the danger of connecting agents to live production databases due to unpredictable agent behaviours.

InfoTrack eliminates this risk through:

A. Controlled, permissioned access

AI agents can only access data they are explicitly authorised to see.

B. Auditability

Every MCP interaction is logged end-to-end, satisfying OAIC, APRA CPG 234 and ISO 42001 requirements.

C. Real-time but safe

Agents interact with replicated, cleansed, controlled datasets, not production systems.

D. Provenance guarantees

Every datapoint is verified, timestamped and traceable back to an official government source.

This is the “governance by design” framework outlined in InfoTrack’s infrastructure blueprint.

InfoTrack

7.

Strategic roadmap for tech leaders (2026–2030)

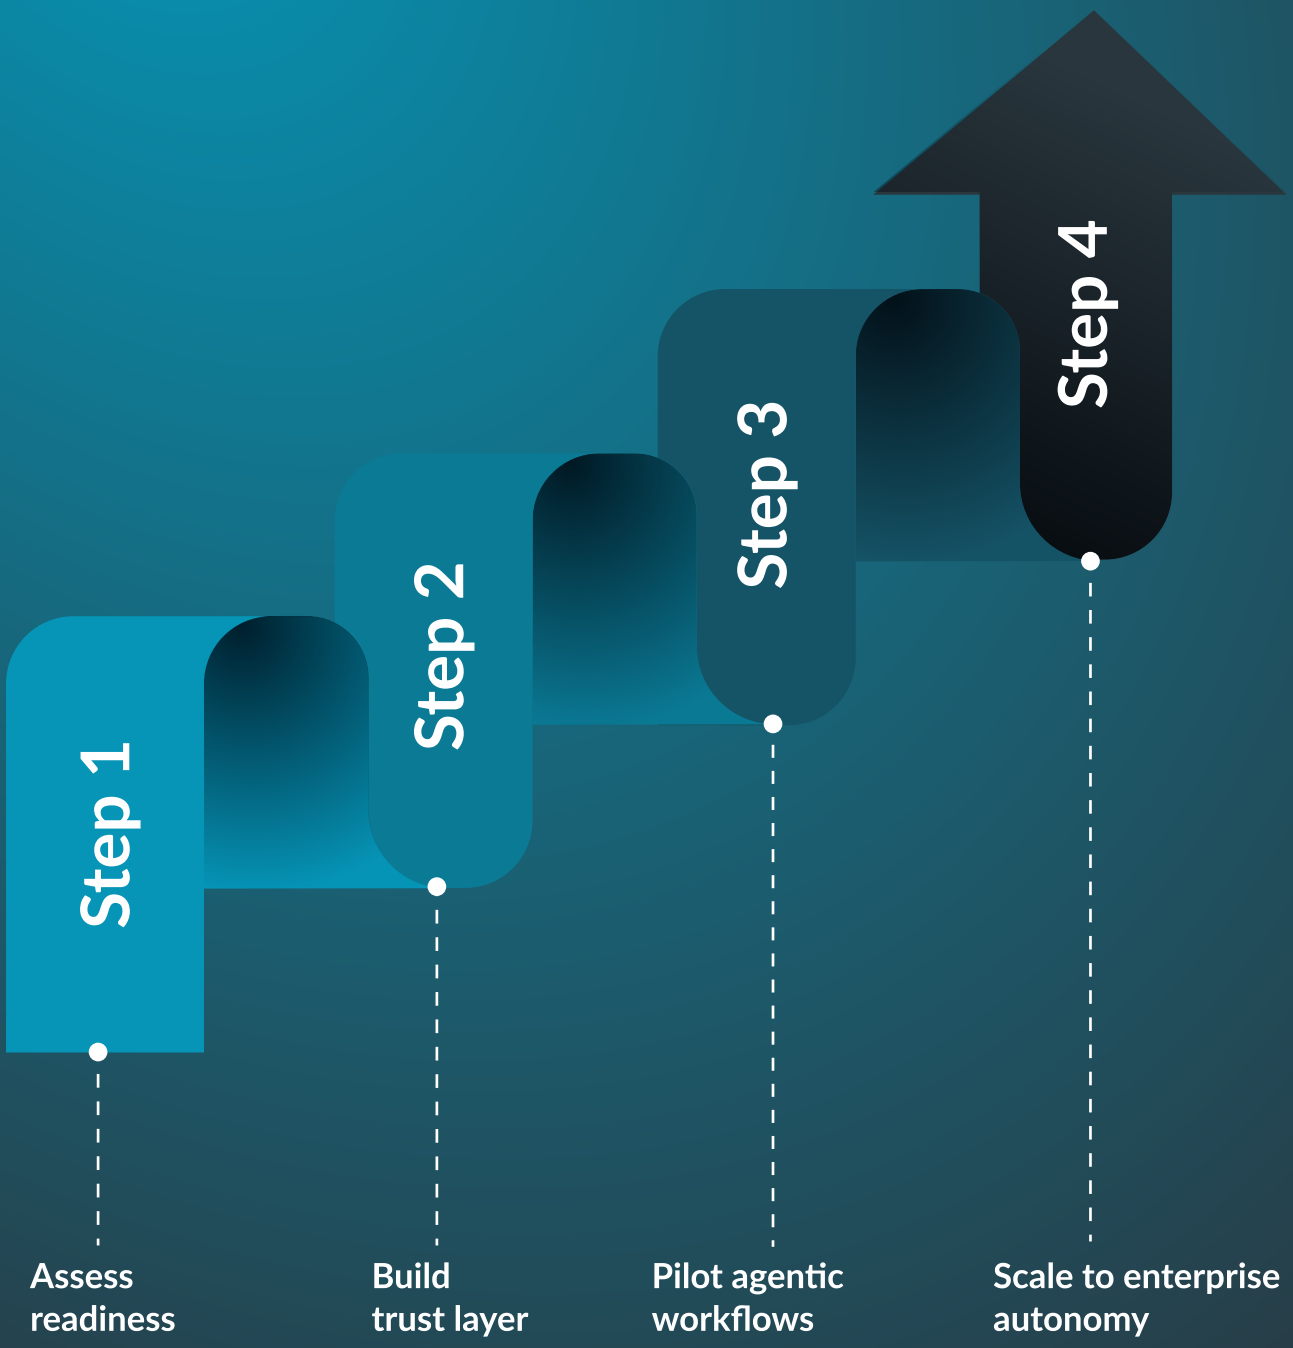

Agentic AI integration workflow (2026–2030)

1. Assess readiness
2. Build trust layer
3. Pilot agentic workflows
4. Scale to enterprise autonomy

Step 1 — Assess readiness (0–90 days)

Evaluate organisational, technical, and compliance maturity to determine where Agentic AI can safely augment high-value workflows.

Step 2 — Build the trust layer (90–180 days)

Deploy secure, governed integration pathways (e.g., InfoTrack MCP Gateway) and ensure identity, auditability, and interoperability foundations are in place.

Step 3 — Pilot agentic workflows (6–12 months)

Test autonomous and semi-autonomous AI capabilities in critical but controlled processes, validating value, compliance, and risk posture.

Step 4 — Scale to enterprise autonomy (12–36 months)

Expand from isolated pilots to organisation-wide adoption migrating to automated, predictive, and event-driven operating models.

Phase 1 — Readiness Assessment (0–90 days)

Focus:
Establish a baseline understanding of AI opportunity, risk, and technical feasibility.

Key activities

- Map high-value, compliance-heavy workflows
  (e.g. identity verification, due diligence, multi-step regulatory submissions)
- Evaluate data lineage and trust boundaries
  Ensure data sources, transformations, and custodianship can be tracked end-to-end.
- Assess AI-readiness of core enterprise systems
  Look at APIs, security posture, interoperability, content schemas, and cloud maturity.
- Identify low-risk, high-impact pilot opportunities
  Select 1–3 workflows where automation improves accuracy, speed, or compliance.

Outcome
A prioritised roadmap of pilot-ready workflows and the technical prerequisites to activate them.

Phase 2 — Establish the Trust Layer (90–180 days)

Focus:
Build the secure, auditable connective tissue between AI agents, enterprise systems, and government services.

Key activities

- Deploy the InfoTrack MCP Gateway
  Creates a compliant, pre-audited pipeline for AI-to-government connectivity.
- Connect AI agent frameworks
  e.g. Microsoft Copilot, OpenAI Agents, Salesforce Einstein, Google Vertex AI.
- Implement role-based access controls (RBAC)
  Ensure agents only act within approved boundaries.
- Configure audit streams
  Generate real-time logs of agent actions, decision pathways, and data exchanges.

Outcome
A governed, secure, enterprise-grade foundation for Agentic AI execution.

Phase 3 — Pilot Agentic Workflows (6–12 months)

Focus:
Validate real-world performance of autonomous workflows under controlled conditions.

Example pilot use cases

- Autonomous identity verification
  AI agents interpret documents, lodge checks, retrieve verification results.
- Automated government data retrieval
  Agents initiate searches, request documents, and populate case files.
- AI-driven compliance reporting & filings
  Auto-generate reports, cross-validate data, submit filings on behalf of humans.
- Real-time case, customer, or asset intelligence
  Agents continuously update dashboards, alerts, and risk signals.

Outcome
Evidence of accuracy, reliability, ROI, and regulatory alignment, enabling scale-up decisions.

Phase 4 — Enterprise-wide Autonomy (12–36 months)

Focus:
Transform from process-by-process optimisation to a fully autonomous digital operating model.

Evolution path

Move from:

- **Manual → Automated**
  Routine tasks performed independently by agents.
- **Reactive → Predictive**
  Agents anticipate needs, pre-fetch data, pre-validate risks.
- **Disconnected → Event-driven**
  Systems respond in real time to trigger conditions, eliminating process bottlenecks.

Outcome
A resilient, AI-enabled enterprise that operates with continuous intelligence, reduced friction, and improved compliance posture.

8.

Australia's future agentic AI ecosystem

From experimentation to national infrastructure, Australia's AI landscape is evolving fast:

1. **Agentic AI becomes foundational**
   Moving from individual productivity to organisational and sovereign productivity.

2. **MCP emerges as the national standard**
   MCP is rapidly becoming the interoperability spine of modern AI architectures.

3. **Trust becomes the competitive differentiator**
   Winners will be those who can prove compliance at scale.

4. **Enterprises shift from tool adoption to system integration**
   Strategy, not experiments, delivers value.

5. **AI-to-government connectivity becomes an economic enabler**
   InfoTrack is positioned to set the gold standard for that trust boundary.

9.

Conclusion

Agentic AI represents the most immense opportunity in enterprise technology since the cloud.

But unlike the cloud's infrastructure shift, this is a foundational leap to autonomy and it comes with a crucial requirement: trust in every action.

InfoTrack's MCP Platform provides the missing trust layer that allows Australian enterprises to:

- Use AI agents safely
- Access government data compliantly
- Operationalise AI workflows with confidence
- Innovate without compromising security

As adoption accelerates, the question for leaders in technology is no longer if to deploy agentic AI but how to deploy it responsibly, securely, and at scale.

About InfoTrack

InfoTrack is Australia's leading provider of secure digital connectivity to government systems. Our MCP platform extends two decades of trusted infrastructure into the era of agentic AI providing compliant, auditable and real-time connectivity between enterprises, AI agents, and verified government data sources.

For invitations to the CIO/CTO Roundtable Series or to explore pilot opportunities:

Visit InfoTrack.ai