

Overseeing vs. Overlooking AI:

Are You Really Confident You
Know What's Going On?

Whitepaper

 Digital Service
Management




Overseeing vs. Overlooking AI: Are You Really Confident You Know What's Going On?

Across the enterprise and government sectors, a widening AI gap has formed between executive rhetoric and operational reality. While organisations adopt AI to signal innovation, they often do so without defined goals or the data foundations required to achieve them.

Consequently, many who believe they are overseeing AI adoption are actually overlooking its full potential and operational risk. To scale AI effectively, organisations must move beyond binary monitoring to a governance system built on deep observability — one that provides the accountability, explainability and data integrity needed for a genuine competitive advantage.

You can't manage what you can't see, and right now, many are flying blind.





The Integration Deficit

Paying for Potential, Not Performance

History is repeating itself: AI adoption is mirroring the early days of cloud migration. The race to demonstrate innovation with artificial intelligence is driving significant cost spikes long before value is realised. With major vendors like Microsoft, Apple, Salesforce and ServiceNow embedding AI directly into their licensing models, organisations are now absorbing the cost of these capabilities before they've even built a use case.

Beyond mature examples (such as chatbots), most AI functionalities, while heavily marketed, remain largely AI-washing. Without extensive configuration and integration, these tools cannot deliver outcomes.

As Rick Mathot, Group Director at Versent, explains,

"This is the 'AI Gap': the gap between potential versus implemented, and aspirations versus commitment."

And that disconnect is why efficiency gains aren't hitting the anticipated 20 to 30% mark; Deloitte's 2025 survey shows only [6% of executives saw a payback](#) within a year.

By overestimating data maturity and ignoring infrastructure gaps, organisations are paying for siloed tech islands that fail to produce results.



Shifting From Monitoring to Observability

Bridging this gap requires a shift from passive monitoring to active observability. Traditional monitoring practices, even when driven by AI, typically provide binary data indicating whether a system is functional or broken. It focuses on what happened rather than the broader context of why it happened.

Observability, by contrast, correlates data across the entire operational landscape to provide transparency and accountability. It explains *why* and *how* events occur, justifying a system's decisions against business objectives.

This depth of insight is a prerequisite for autonomous agentic decision-making. To be effective, governance must move beyond historically reactive diagnostics to address:

- **Accountability and explainability:** Ensure every autonomous decision is defensible. If a system cannot explain its logic, it should not act.
- **Value realisation:** Moving beyond ticket closure rates to demonstrate AI is delivering specific efficiencies, based on clearly defined objectives for AI outcomes.
- **Risk management:** Monitoring for drift (when an AI model's behaviour or outputs deviate over time) and hallucinations (when AI generates incorrect, nonsensical or unverified information) to ensure autonomous agents operate within security posture and intended business logic.


Only by maturing monitoring into observability can organisations enable outcomes beyond mere incident prevention. This transition relies on a continuous feedback loop in which governance and observability are inextricably linked, mutually reinforcing the transparency, accountability and direction required for true digital transformation in the AI age.



Contrary to current adoption trends, this shift is not achievable by simply deploying a plug-and-play AI capability. Dr Tony Parsons, Head of Digital Service Management & Chief Strategist at Versent, advises,

"If you don't have a consolidated data foundation for AI to perform analytics, derivation and instantiation, AI just doesn't work."

To achieve meaningful outcomes, organisations must rethink data capture and consolidation across functions and tools. That calls for a deeper restructuring of operational processes before the gap becomes too wide to close.



Bridging the Gap Between Aspiration and Architecture



If you don't know what you've got, you can't manage it. While the transition from monitoring to observability provides a technical roadmap, market realities will make maturation an operational mandate. Friction still persists because the rush to adopt AI has bypassed the fundamental, boring work of data cleansing and normalisation that make these systems viable.

The momentum toward this new governance model is driven by three factors:

The Rise of the Digital Employee

We are moving toward Agentic AI. This is a shift McKinsey's 2025 research confirms, noting that 23% of organisations are [already scaling agentic systems](#), while another 39% are in active experimentation. Digital agents are now operating as part of the billable workforce. This evolution pushes humans up the loop to handle strategy and more complex technological challenges, while autonomous agents manage the execution volume.

The Shift to Technology Arbitrage

The market is shifting from labour arbitrage to technology arbitrage. As long as organisations absorb the cost of AI features regardless of realised value, avoiding adoption is futile, especially when competitors leverage these tools to lower costs and accelerate time-to-market.

Operational Saturation

Humans increasingly cannot manage the volume of events without AI automation.

"Modern IT environments are far too complex and far too mission-critical to allow for the old 'eyes-on-glass' monitoring," Tony explains.

In sectors like government, the only way to comply with strict security controls while managing this complexity is to adopt contained AI solutions that minimise risk while delivering necessary scale.

The Cost of Static Governance

Bridging the AI gap creates a level of operational agility that static governance can't



match. When approved channels are inefficient, business units turn to public tools, triggering Shadow AI and data sprawl (the proliferation of data and AI processes, resulting in a loss of central data governance oversight).

Without observability to anchor these new technologies, the gap between potential and performance presents measurable risk across an organisation.

Financial and Operational Sustainability

Organisations are already absorbing the cost of digital agents; failing to deploy them means accepting rising overheads while competitors use observability to resolve incidents autonomously. This creates a divide between the cost of operating and the ability to act, which will eventually make pre-AI service models unsustainable.

Technical Integrity and Workforce Agility

The internal dissonance created by siloed AI usage prevents AI from accessing the consolidated data it needs, leading to failed automation and workforce de-skilling. If humans lose deep system knowledge to a process that lacks accountability and explainability, the organisation becomes paralysed when AI encounters complex challenges it was never programmed to resolve.

Security and Governance

Traditional Change Advisory Board models are fundamentally incompatible with the instantaneous nature of AI. Their comparatively glacial pace will eventually force organisations to choose between lagging innovation and losing control of the tech environment. Without policy-as-code operations and deep observability, organisations can't detect when models deviate from their original baseline.


Further, as models become more sophisticated, they're not necessarily hallucinating less but are getting better at *masking* errors. A lack of traceable decision-making allows oversights to scale, meaning technical glitches carry far more serious potential



consequences.

Business Outcomes and Market Velocity

A disconnect between IT performance and revenue poses risks to business operations. Given the availability of instantaneous digital services, an unobserved technical outage can cost millions of dollars in lost revenue per hour. Unless AI performance is tied to specific business outcomes, organisations risk half-baked implementations that increase complexity without delivering the promised velocity.



Failed Automation

Siloed AI usage prevents consolidated data access



Performance & Risk Gap

Unanchored technology creates measurable organisational risk



Unsustainable Models

Rising overheads make pre-AI service models non-competitive



The Roadmap to AI Maturity

From Fragmentation to Autonomous Systems

To progress from paying for tools to realising outcomes, organisations typically move through four stages of adoption. Each phase reveals challenges that mature observability and integrated governance can resolve.

Step 1: Experimentation and Isolation

Shadow AI is the driver of experimentation, in which individual teams adopt disparate tools, such as public LLMs, to address workflow bottlenecks. This stage often bypasses IT oversight.

The technical reality is fragmented, as disjointed operations lack connection to a broader strategy. Data sprawl poses the greatest risk, as proprietary client data leaks into public domains, circumventing centralised governance.

Step 2: Integration and Foundational Trust

Organisations attempt to bridge these silos to drive innovation. However, without a standardised framework, they might find a spaghetti of loose connections in their workflows and outcomes.

For instance, in the public sector, compliance itself can be a wall. Static security protocols (like IRAP) block dynamic tools, creating a standoff between innovation and safety. While auditable, governed data processes establish the trust needed to bridge into the next phase, limited observability creates new challenges. Without it, model drift and hallucinations can go undetected, turning what should be a leap in innovation into an exercise in blind trust.



Step 3: Automation and Decision Velocity

Maturity accelerates as organisations move toward Agentic AI models. These should be supported by pre-approved governance standards. Basic decision-making capabilities transition to the system, allowing human talent to handle higher-level tasks.

However, organisations will likely need to take a leap in data and governance processes leading into this phase, as AI cannot function effectively at this speed. The organisation will likely hit a maturity “glass ceiling” if the supporting data is compromised or incomplete, Tony explains.

Step 4: Adaptive, Autonomous and Converged Systems

The final state is a self-healing hybrid workforce of humans and digital agents. Leaning on predictive analytics, an autonomous system can prevent outages before they occur and explain the comprehensive decision backlog when they do.

For the public sector, this requires contained solutions that minimise risk to enable scale. Rick notes that “business rules as code will evolve into business rules and policy as code.” In this environment, governance is baked directly into the IT Ecosystem.



Outcomes: The Future of AI Governance

Integrating governance and observability transforms blind trust in AI operations into explainable, accountable automation. It positions organisations for an inevitable future where uptime is revenue and predictive analytics provide the speed needed to compete. Ultimately, governance defines the operational boundaries, while observability provides the telemetry to demonstrate that the system operates within them.

As a strategic partner, Versent bridges the gap between the unrealised potential of AI capabilities and the need for value realisation of government and enterprise clients. Our experts provide the integration framework, data normalisation and workflow design, aligned to ITIL Version 5, to make observability and AI governance operational and safe through:

1. **Establishing the integration framework:** implementing an integration framework that normalises and standardises data, allowing AI to access a consolidated, end-to-end view of the environment.
2. **Conducting discovery and gap analysis:** Versent identifies the gaps between an organisation's current state and its operational aspirations, enabling targeted investments that achieve tangible ROI.
3. **Data hygiene and knowledge curation:** Cleaning, correlating and validating data while establishing systems to capture and maintain business-specific knowledge, enabling AI tools to access reliable information, reducing the risk of hallucinations, drift and sprawl.
4. **Contained solutions for risk management:** Domain-specific, certified environments that allow agencies to leverage AI while containing risks to acceptable levels. Autonomous action is fully explainable and auditable.
5. **Organisational and workflow redesign:** Frameworks and use of Organisational Change Management to update human operating models, establish new roles and adjust internal decision-making velocities to match the speed of AI.



Building the Foundation for Practical AI Performance

The transition from overlooking AI's capabilities to overseeing them requires a shift in governance, operational workflows and data integrity. When governance is the catalyst for and consequence of comprehensive observability, government and enterprises can achieve accountable, explainable and automated scalability. Only then can organisations realise genuine value from the technology and close the gap between aspirations and their current reality.

The process starts with data normalisation and embedding governance frameworks early, rather than retrofitting them down the road.

Partner with Versent to align your data foundations and transform AI into a controlled, observable engine for growth.

Contact us →

Versent Pty Limited

Level 11/111 Elizabeth St,
Sydney NSW 2000

enquiries@versent.com.au
(03) 8680 2363


**Digital Service
Management**