idm.
information & data manager

APRIL-MAY 2026

AI Chatbots Must Earn Citizen Trust

AI Safety Takes Backseat in National Plan

Nine Reasons Failing AI Projects Never Die

Budget Cuts Hit Privacy Teams Hard

FIIG Hit with $2.5m Cyber Security Penalty

Australian fixed-income specialist FIIG Securities has been ordered to pay $2.5 million in penalties after cyber security failures exposed 18,000 clients to a data breach that saw 385 gigabytes of confidential information stolen.

The Federal Court penalty marks the first time civil penalties have been imposed for cyber security failures under general Australian Financial Services licence obligations, setting new compliance expectations for the financial services sector.

ASIC brought the case against FIIG Securities Limited for failing to protect clients from cyber security threats between March 2019 and June 2023.

The company's inadequate controls worsened a 2023 cyber-attack that leaked driver's licences, passport information, bank account details and tax file numbers onto the dark web.

FIIG admitted it failed to comply with its AFS licence obligations and that adequate cyber security measures would have enabled earlier detection and response. The company also admitted complying with its own policies could have prevented some or all client information from being downloaded.

The Court ordered FIIG to pay $500,000 towards ASIC's costs and undertake a compliance programme involving an independent expert approved by ASIC to ensure its cyber security and cyber resilience systems are reasonably managed.

FIIG's cyber security failures included not allocating necessary financial resources for suitably qualified personnel or adequate technological resources. The company did not implement multi-factor authentication for remote access users, strong passwords and access controls for privileged accounts, or appropriate firewall and security software configuration.

The firm also lacked regular penetration testing and vulnerability scanning, had no structured plan for software security updates, no qualified IT personnel monitoring threat alerts, and no mandatory cyber security awareness training for staff. FIIG did not have an appropriate cyber incident response plan tested at least annually.

"Cyber-attacks and data breaches are escalating in both scale and sophistication, and inadequate controls put clients and companies at real risk," ASIC Deputy Chair Sarah Court said.

"ASIC expects financial services licensees to be on the front foot every day to protect their clients. FIIG wasn't - and they put thousands of clients at risk."

Court noted the consequences far exceeded what it would have cost FIIG to implement adequate controls initially.

"This is the first time the Federal Court has imposed civil penalties for cyber security failures under the general AFS licensee obligations, setting a clear licence-to-operate expectation for robust cyber resilience," she said.

Law firm Herbert Smith Freehills Kramer recommends that "organisations review their cybersecurity settings against the cybersecurity measures steps that were agreed to be "adequate" in this case to consider their appropriateness in their setting."

OpenText Sheds Vertica in AI Pivot

OpenText has agreed to sell its Vertica analytics database platform to Rocket Software for $ US 150 million cash, marking the Canadian software giant's second major divestiture to the same buyer in less than two years.

The transaction, subject to regulatory approval, is expected to close during OpenText's fiscal year 2026. OpenText said it intends to use proceeds to reduce outstanding debt.

The company's total debt load following its 2023 acquisition of Micro Focus for approximately $US6 billion has been a recurring concern for investors.

Vertica is a high-performance, enterprise-grade analytics database designed to handle petabyte-scale data workloads at speed. Unlike conventional transactional databases, Vertica uses columnar storage - organising data by column rather than row - enabling significantly faster analytical queries across large datasets.

It supports massively parallel processing, distributing queries across multiple servers simultaneously, and allows data scientists to run machine-learning models directly inside the database without moving data to external tools.

Under the agreement, Vertica's software, customer contracts, associated services and employees will transfer to Rocket Software, a portfolio company of private equity firm Bain Capital.

idm.
information & data manager

Publisher/Editor: Bill Dawes

Email: bill@idm.net.au

Web Development & Maintenance: Cordelta

Advertising Phone: 02 90432943

Email: idm@idm.net.au

Published by Transmit Media Pty Ltd

PO Box 392, Paddington NSW 2021, Australia

Shadow AI Agents Pose Security Risk

Nearly one-third of employees use unsanctioned AI agents for work tasks, creating security vulnerabilities most organisations cannot address, a new Microsoft report reveals.

The Cyber Pulse AI Security Report, reveals 29% of employees use unauthorised AI agents. Only 47% of organisations implement specific generative AI security controls.

Microsoft's research shows over 80% of Fortune 500 companies now deploy active agents built with low-code and no-code tools. However, rapid deployment is outpacing security and compliance frameworks, creating what Microsoft terms "shadow AI" risks.

"The dual nature of AI has arrived: extraordinary innovation paired with unprecedented risks," the report states.

Microsoft's Defender team recently identified a fraudulent campaign exploiting "memory poisoning" - an attack technique that manipulates AI assistants' memory to persistently steer responses. The company's AI Red Team documented how agents were misled by deceptive interface elements embedded in everyday content.

"We need to treat agents like humans and apply Zero Trust principles," said Vasu Jakkal, corporate vice president at Microsoft Security.

The report highlights that AI agent adoption spans all regions and industries globally. Financial services represents 11% of active agents worldwide, manufacturing accounts for 13%, and retail comprises 9% of global agent usage.

Microsoft warns that unsupervised or ungoverned AI agents threaten security, business continuity and reputation. Agents with excessive access or incorrect instructions become vulnerabilities. Bad actors can exploit these as "double agents."

The company recommends organisations establish observability through centralised registries, identity-based access controls and real-time monitoring dashboards. This includes cross-platform governance and built-in security protections.

Microsoft's seven-point governance checklist includes documenting each agent's purpose with least-privilege access. It also recommends applying data protection rules to AI channels and offering secure alternatives to curb shadow AI. Organisations should update business continuity playbooks and elevate AI risk to board-level visibility.

The report introduces Agent 365, Microsoft's unified control plane for managing AI agents across organisations. The platform provides centralised registration, governance, security observation and operation for agents built on Microsoft platforms, open-source frameworks or third-party systems.

Microsoft emphasises that organisations succeeding with AI agents prioritise observability, governance and security. This requires collaboration across IT, security, AI teams and developers through unified control platforms.

Westpac Rolls Out Copilot Worldwide

Westpac is deploying a global rollout of Microsoft 365 Copilot, giving artificial intelligence tools to its entire workforce of 35,000 employees plus contractors.

The deployment, which Westpac claims is the largest in financial services within Asia-Pacific, follows a pilot programme involving 15,000 employees in Australia. The bank has not disclosed the financial investment or specific productivity metrics from the trial.

The rollout includes Microsoft's Copilot Studio, which allows IT teams to build custom AI agents. Westpac said it has developed agents for HR and IT functions to handle employee queries and routine tasks, though the bank did not specify how many custom agents are operational.

Andrew McMullan, Westpac's Chief Data, Digital and AI Officer, said AI needs to be paired with governance frameworks.

"Technology alone isn't the answer," McMullan said. "We know AI needs to be used responsibly and has the strongest impact when it's paired with skilled people, strong values and good judgement."

The bank is providing AI training through masterclasses and prompt-a-thon workshops.

It has also established an innovation sandbox on Microsoft Azure for experimentation with AI-driven solutions.

The move reflects broader adoption of generative AI across Australian financial institutions, which must balance productivity gains with regulatory compliance requirements.

Banking regulators globally have emphasised the need for governance frameworks around AI systems that process customer data.

Nine Reasons Failing AI Projects Never Die

If you have ever suspected that failure is not an obstacle for generative AI projects, new research from Finland may help explain why.

A peer-reviewed study published in Big Data & Society this month tracked the development of a large language model (LLM)-based decision support tool inside a Finnish public sector organisation over nearly 18 months. The tool was designed to help welfare claims specialists navigate complex, scattered and frequently updated guidance documents - a challenge familiar to compliance, records and information managers worldwide.

Despite repeated failures to deliver accurate, precise and consistent outputs across five rounds of user testing, the project never stopped. The researchers identified nine distinct "justification frames" that made the tool - in their words - "an irresistible thing to influential organizational actors."

The tool functioned like a ChatGPT-style bot. It searched a database of guideline documents and returned answers with links to source material. But experiment after experiment showed it could not reliably meet user expectations. One lead innovator reportedly cited tool accuracy at "75 to 80 per cent" throughout the project. When a manager pressed for verification, the response was: "75% is my recollection. But I can check."

Nine Frames That Shield a Failing Project

Researchers Marta Choroszewicz of the University of Eastern Finland and Antti Rannisto of Aalto University categorised the justifications into three groups: tool-oriented, process-oriented and ideology-oriented.

The five tool-oriented frames drew on promises of efficiency, cost savings, employee well-being, fairness for citizens and organisational desirability. When the fourth experiment produced decidedly mixed results, internal narratives nonetheless circulated that the tool was performing well. Other business units began requesting access.

"Everyone is really excited about this [tool]!" an innovator told internal meetings. "The attraction around it is huge."

The market frame - cost savings - proved impossible to substantiate. Experimental results suggested that realising the tool's value would require excessive additional resources from frontline workers.

Yet, as the paper records, when a team member asked whether the business impact had been calculated, the leading innovator replied: "No, but the business impact of this is, like, really big."

The civic frame proved particularly powerful in the Nordic welfare state context. Innovators framed the tool as a pathway to fairer, more consistent decisions for citizens. As one lead innovator stated in media appearances: "It is impossible to respond to changing individual needs without smart automation or AI."

The remaining four frames addressed the process itself. Rapid, agile experimentation - borrowed from private sector design industry practice - was presented as evidence of organisational maturity. Failures were reframed as necessary learning. "The important thing here would be that 'continuous development' - it is already delivering value and benefits," a lead innovator told the team. "It already deserves to exist."

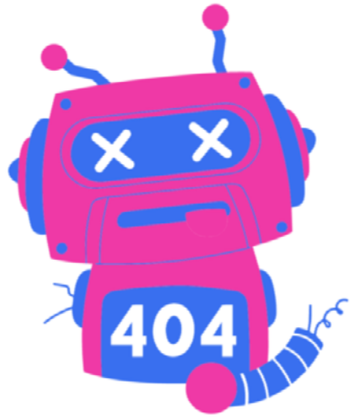

Blame the User, Not the Tool

When the tool's outputs proved unreliable, innovators shifted focus to users and organisational processes rather than the technology itself. "The challenge lies in users' inability to effectively interact with it - they are tied too much to the old way of working," innovators and consultants argued across multiple internal meetings.

The researchers describe this dynamic as "the constructivist aspect of technosolutionism" - redefining the problem to fit the pre-existing solution, shifting the burden of adaptation onto users and organisations rather than questioning the tool's fitness for purpose.

Choroszewicz notes: "Our findings show that generative AI projects in public administration often continue not because the tools work well, but because a compelling set of justifications makes them hard to stop and declare the tools non-functional."

Co-author Rannisto adds: "By normalizing setbacks as learning, the team maintained innovation momentum even when accuracy, precision and consistency remained out of reach."

Why GenAI Is Particularly Hard to Challenge

The researchers argue that LLM-based tools present a specific governance challenge compared to older, more deterministic AI systems. Because they function as "black boxes," innovators cannot easily identify why they fail or how to fix them. This opacity drives attention away from technical problems and toward changing users and workflows.

"Instead of treating generative AI tools' failures as chances for substantive learning, reconsidering technical alternatives or pausing development, failures are often reframed as expected steps in a 'learning process,'" the authors write.

A three-frame combination proved especially durable across the organisation. The industrial frame (efficiency), vitalist frame (worker well-being) and civic frame (equal treatment of citizens) formed what the researchers call "a powerful trio" capable of crossing organisational and professional boundaries - particularly resonant in welfare state contexts where public value, staff welfare and service equity all matter.

The paper is available at Big Data & Society: https://doi.org/10.1177/20539517261424159

Legacy IT Still Blocking Government Cyber Uplift

Australian government entities have achieved strong protective security compliance results in 2024-25, but critical gaps in technology and cyber maturity remain unresolved, two landmark reports show.

The Department of Home Affairs' Protective Security Policy Framework (PSPF) Assessment Report 2024-25 and the Australian Signals Directorate's (ASD) Commonwealth Cyber Security Posture paint a detailed picture of where Australian government security stands - and where it falls short.

The PSPF report covers 100 Non-Corporate Commonwealth Entities (NCEs) reporting under a new compliance-based model introduced in November 2024.

The ASD report is the sixth annual cyber security posture report tabled before Parliament.

The PSPF Assessment Report found 92 per cent of entities achieved an overall "Effective compliance" rating. No entity recorded an overall "Low compliance" result.

However, the Technology domain - covering ICT lifecycle management, cyber security strategies and programs - was the clear weak point. Just 79 per cent of entities achieved Effective compliance in Technology, the lowest of all six security domains. Seventeen per cent reported Moderate compliance and four per cent Low compliance.

The report notes the Technology domain result "is consistent with previous reporting periods," suggesting this gap is structural rather than incidental.

By contrast, the Information domain - covering classification, information holdings, disposal and sharing - recorded the highest Effective compliance rate at 96 per cent.

PSPF Directions: Mandatory Actions Completed

Five mandatory PSPF Directions were issued during 2024-25, covering Foreign Ownership Control or Influence (FOCI) risks, technology asset stocktakes, ASD cyber security partnership engagement, and the prohibition of DeepSeek and Kaspersky products on government systems.

All entities reported full compliance with all five Directions.

The prohibition on DeepSeek products, covered under Direction 001-2025, follows growing international concern about Chinese-developed AI applications accessing government systems and data. Several allied governments including the United States, United Kingdom and European Union institutions moved to restrict or ban DeepSeek on government devices in early 2025.

The ban on Kaspersky Lab products under Direction 002-2025 reflects longstanding concerns about Russian-linked software within critical government infrastructure.

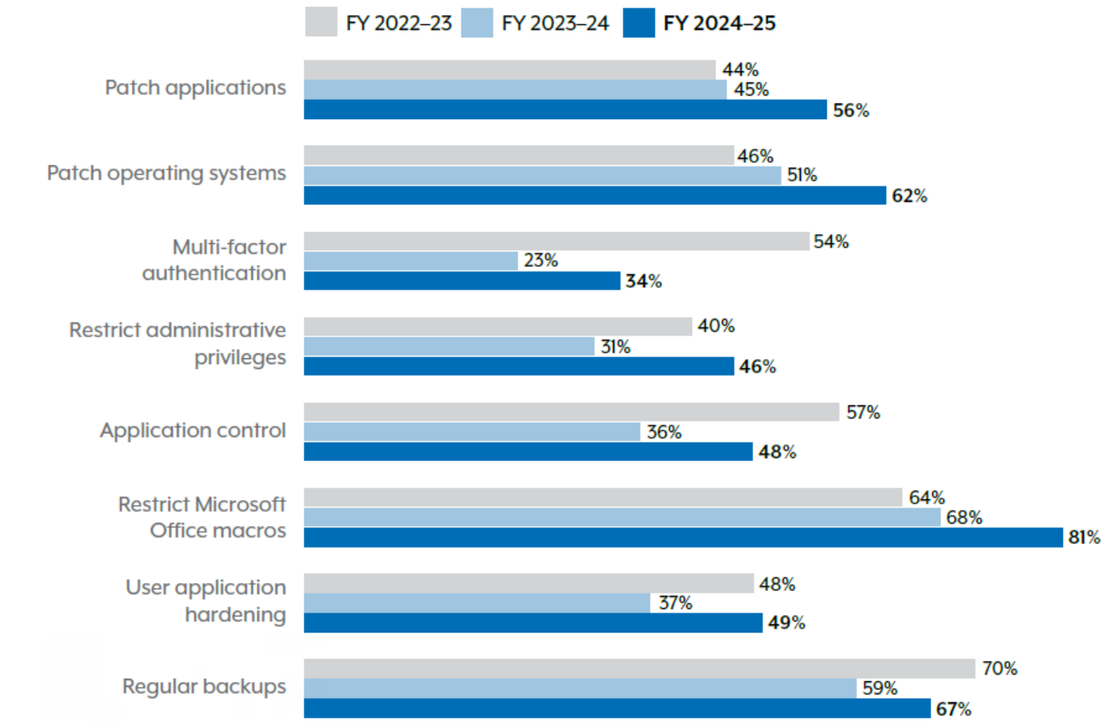

Government Entities that reached Essential Eight Maturity Level 2 or higher.

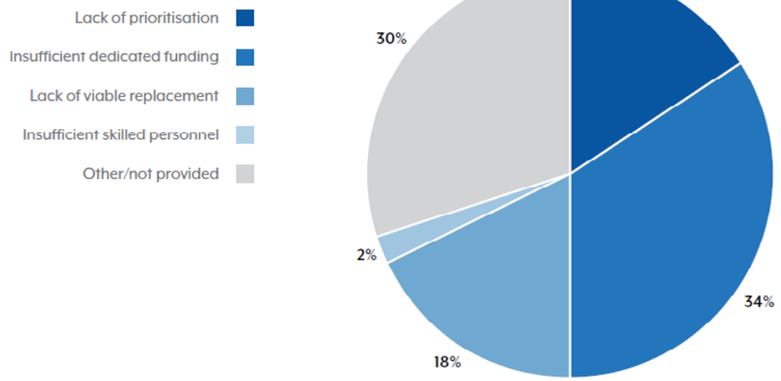

The most significant reasons for using legacy IT

Top Security Risks Identified

The top five security risks identified by entities in 2024-25 were: compromise or unauthorised disclosure (58 per cent); trusted or malicious insider (58 per cent); cyber security emerging threats (55 per cent); funding, resources or capability limitations (49 per cent); and cyber security attack (43 per cent).

Governance Risk and Compliance Managers will note that insider threat - both trusted and malicious - ranked equal first alongside data compromise.

The report states: "These identified threats are consistent with previous reporting periods."

Essential Eight Maturity Remains a Challenge

The ASD Cyber Security Posture report highlights that achieving Essential Eight Maturity Level 2 - mandated for all NCEs since 1 July 2022 - remains elusive for most entities.

In 2024-25, just 22 per cent of all government entities achieved Maturity Level 2 or higher across all eight mitigation strategies. This was up from 15 per cent in 2024, but still well below the 25 per cent recorded in 2023.

The report explains the 2023 decline was caused by ASD hardening the Maturity Level 2 controls in November 2023. Key changes included stricter phishing-resistant multi-factor authentication (MFA) requirements, updated application control rules and a revised approach to data backup prioritisation based on business criticality.

The report identifies the most problematic strategy: MFA reached Maturity Level 2 in only 34 per cent of entities in FY2024-25, though this is up from 23 per cent the prior year.

Restrict Microsoft Office macros was the strongest performer at 81 per cent.

Legacy IT: Persistent Barrier to Uplift

Legacy IT continues to impede cyber security improvements across government. In 2025, 59 per cent of entities said legacy technology impacted their ability to implement the Essential Eight - down from 71 per cent in 2024, but still representing the majority.

Entities reported insufficient dedicated funding (34 per cent) and lack of a viable replacement (30 per cent) as the most significant reasons for continued legacy IT use.

Despite entities reporting more incidents internally - with 62 per cent reporting at least 80 per cent of incidents to senior executives - external reporting to ASD remains low.

Only 35 per cent of entities reported at least half of observed cyber security incidents to ASD in 2024-25. This is up from 32 per cent in 2023-24 but well below what ASD considers adequate.

The report notes: "Any degradation in the quantity or quality of information reported to ASD reduces our capacity to support the entity to mitigate the impacts of cyber compromise."

ASD responded to 408 cyber security incidents from government entities in 2024-25, representing 33 per cent of all incidents responded to nationally.

Leadership and planning indicators showed improvement. Eighty-two per cent of entities had a documented cyber security strategy in 2025, up from 75 per cent in 2024.

Cyber security training improved significantly, with 87 per cent of entities providing annual training to their full workforce, up from 78 per cent in 2024.

However, privileged user training - critical for IT administrators and system owners - declined to 45 per cent in 2025, down from 51 per cent in 2024.

Supply chain risk assessments also declined, with 70 per cent of entities performing them in 2025, down from 74 per cent in 2024. This is a concern for procurement managers and enterprise architects responsible for vendor risk management.

Post-Quantum Cryptography: A New Priority

The ASD report flags an emerging compliance obligation for IT and security teams. All organisations are encouraged to begin transitioning to post-quantum cryptography by 2030. ASD released new cryptography guidelines in September 2025.

The report warns: "CRQC will render common public-key encryption protocols insecure. This means communications, information and data once thought secure, could be at a greater risk of compromise."

A new PSPF Assurance Capability introduced under the 2023-2030 Australian Cyber Security Strategy will pilot verification of entity self-assessments against the 2024-25 results.

This is designed to address what the report calls "optimism bias commonly associated with self-assessment."

The PSPF Assessment Report is available at www.protectivesecurity.gov.au. The Commonwealth Cyber Security Posture in 2025 is available at www.cyber.gov.au.

AI Creates Workload Creep

Generative AI tools are intensifying workloads rather than reducing them, according to new research from UC Berkeley. Employees are working faster, taking on broader tasks and extending work hours.

An eight-month study of 200 employees at a US technology company found workers voluntarily expanded their responsibilities. AI made "doing more" feel possible and rewarding.

"AI tools didn't reduce work, they consistently intensified it," researchers Aruna Ranganathan and Xingqi Maggie Ye wrote in Harvard Business Review.

The research, conducted from April to December last year, identified three forms of work intensification that organisations implementing AI-driven automation should anticipate.

Workers increasingly stepped into responsibilities that previously belonged to others once AI filled knowledge gaps. Product managers and designers began writing code, while researchers took on engineering tasks.

"Workers increasingly absorbed work that might previously have justified additional help or headcount," the researchers found.

Engineers spent more time reviewing AI-generated work produced by colleagues, including "vibe-coding" attempts and partially complete code. This oversight surfaced informally through Slack threads and desk-side consultations, adding to workloads.

Blurred Work Boundaries

AI reduced the friction of starting tasks, causing workers to slip small amounts of work into breaks. Many prompted AI during lunch, meetings or while waiting for files to load.

"Work felt less bounded and more ambient - something that could always be advanced a little further," the researchers observed.

Workers described realising in hindsight that prompting during breaks became habitual, eliminating recovery time. The conversational style of prompting made work spill into evenings without deliberate intention.

AI introduced a rhythm where workers managed several active threads simultaneously. They manually wrote code while AI generated alternatives, ran multiple agents in parallel or revived deferred tasks.

"You had thought that maybe, oh, because you could be more productive with AI, then you save some time, you can work less," one engineer told researchers.

"But then really, you don't work less. You just work the same amount or even more."

The voluntary expansion of work creates a self-reinforcing cycle where AI acceleration raises speed expectations, making workers more reliant on AI tools.

"What looks like higher productivity in the short run can mask silent workload creep and growing cognitive strain," the researchers warned.

Without intervention, overwork can impair judgement, increase errors and make it harder to distinguish genuine productivity gains from unsustainable intensity. Workers face fatigue, burnout and difficulty stepping away from work.

Building an 'AI Practice'

Organisations should develop intentional norms and routines to structure how AI is used. This "AI practice" should define when work should not expand.

Recommendations include intentional pauses before major decisions to assess alignment and reconsider assumptions. Work should be sequenced to advance in coherent phases rather than requiring continuous responsiveness. Organisations should protect time for human connection to counter individualising effects.

"The question facing organisations is not whether AI will change work," the researchers concluded. "But whether they will actively shape that change - or let it quietly shape them."

The study used in-person observation two days weekly and tracked internal communication channels. Researchers conducted more than 40 in-depth interviews across engineering, product, design, research and operations.

The company did not mandate AI use but offered enterprise subscriptions to commercially available AI tools.

Blacktown Council deploys AI assistant for Development Applications

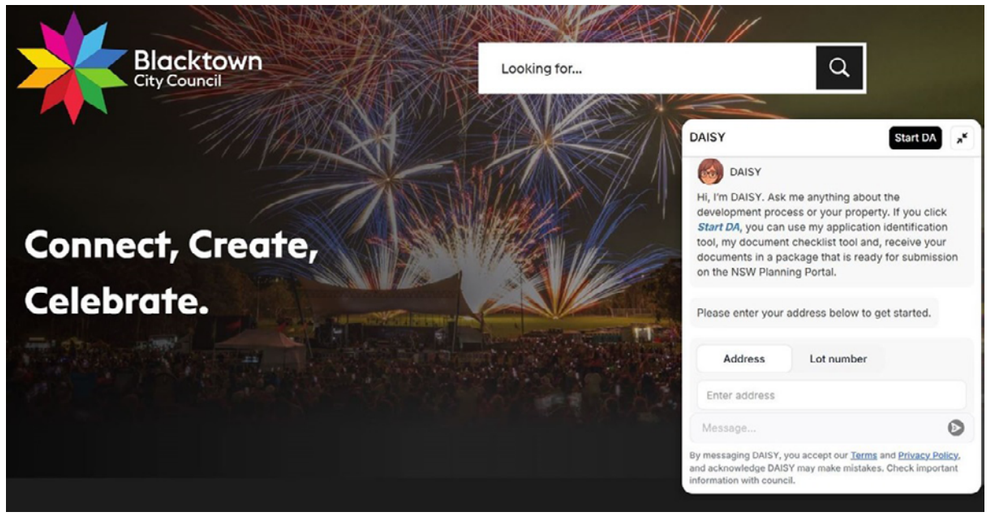

Blacktown City Council has become one of 16 NSW councils deploying artificial intelligence tools to handle planning enquiries. DAISY (Development Application Information System) is an AI-powered digital assistant that provides 24-hour access to planning information and guidance.

The system aims to improve application quality at the point of submission, potentially reducing processing delays.

More than 80 per cent of development applications received by the council come from individual residents seeking to build homes, renovate properties, or subdivide for dual occupancy. The council approved 2,248 residential dwellings in 2023-24.

DAISY helps users understand planning controls, identify required documentation, and check basic submission requirements before lodging applications through the NSW Planning Portal.

Mayor Brad Bunting said the tool was introduced to make the planning process clearer and more accessible for the community.

"For many residents, the development application process can feel complex and overwhelming," Mayor Bunting said. "DAISY is designed to give people clearer information earlier, so they can submit better-prepared applications and avoid unnecessary delays."

He emphasised DAISY does not replace council planning staff or decision-making processes, describing it as a practical customer service tool.

The implementation forms part of the NSW Government's broader push to accelerate housing approvals through AI automation. In September 2025, Planning Minister Paul Scully announced NSW would integrate AI into State Significant Development assessments, positioning the state as "the only state in Australia adopting AI for larger-scale developments".

The NSW Government's Artificial Intelligence in NSW Planning - Council Early Adopter Grant Program provided funding for DAISY, though the grant amount has not been disclosed. The program, administered by the Department of Planning, Housing and Infrastructure, distributed more than $2.7 million to 16 councils to trial AI-based planning solutions.

Blacktown partnered with Adaptovate, a Sydney-based transformation consultancy founded in 2017, to develop the tool. Adaptovate specialises in agile transformation programs and AI implementation across government and private sector organisations.

The grant program required successful councils to report on progress throughout the project, with completion mandated by 30 June 2025.

Several Victorian councils have deployed similar AI planning assistants. Yarra Ranges Council, Mornington Peninsula Shire Council, and Bayside City Council implemented myLot, a locally developed AI planning enquiry service, with implementation periods of two to three months.

The deployment aligns with national policy initiatives encouraging AI adoption in government services. The Australian Government released its National AI Plan in December 2025, emphasising public sector productivity improvements through AI tools.

Federal Treasurer Jim Chalmers urged state departments to "unlock more housing" through AI automation following the NSW announcement, describing it as essential for productivity growth as Australia pursues its target of 1.2 million new homes by 2029.

The NSW Government's AI Assessment Framework and AI Ethics Policy require transparency in AI deployments, including meaningful information about decision-making processes and data usage. Councils must be able to explain how algorithms operate and protect intellectual property rights when engaging external AI providers.

Blacktown City Council stated it will continue to refine DAISY based on user feedback and ongoing improvements to planning data and systems. The tool is currently available on the council's website.

https://www.blacktown.nsw.gov.au

Emerging Data Sources Create ‘All-Time High’ for Digital Risk: Report

New research reveals 88% of legal departments are apprehensive about risks from cloud, chat and collaboration platforms, with most organizations unprepared for modern data challenges

Data risks and compliance challenges have reached unprecedented levels in 2025, with legal teams increasingly overwhelmed by the complexities of managing information across cloud platforms, chat applications and collaboration tools, according to new research from FTI Consulting.

The comprehensive report, The State of Emerging Data in 2025: Tracking a Decade of Data Challenges, reveals that 88% of legal department leaders are now apprehensive about risks surrounding emerging data sources – a dramatic increase from previous years, with 59% expressing they are “very” or “extremely” concerned compared to just 30% in the prior year.

The research highlights a fundamental shift in how organizations must approach data governance and legal discovery. Nearly all litigation and regulatory investigations are now impacted by challenges in preserving, extracting, analyzing and reviewing data from modern digital platforms.

“Legal professionals are realizing that traditional workflows and case law that were once cornerstones of standard practice can no longer provide default guidelines for digital forensic and discovery methodologies,” the report states.

The scale of the problem is evident in the data: by 2024, roughly 80% of the data processed by FTI Technology for client matters originated from emerging data sources, compared to just 35% in 2022.

Organizations Unprepared for Modern Data Realities

Despite the growing prevalence of these challenges, most organizations remain poorly positioned to address them:

- 65% of organizations are minimally or not at all prepared to handle issues related to emerging data sources
- 47% of legal teams have experienced new governance, compliance and discovery challenges associated with cloud, chat and collaboration tools
- 57% have faced new compliance challenges related to their expanding data footprint

The report identifies ten critical areas where traditional approaches are failing:

1. Defining a document - Chat messages, linked content and collaborative files have fundamentally changed what constitutes a “document” in legal contexts.

2. Shared access - Multiple authors, editors and viewers across cloud platforms have blurred traditional concepts of document custodianship.

3. Chat messaging - Short-form, asynchronous communications with emojis, reactions and casual language present new challenges for legal review and context interpretation.

4. Linked content - Hyperlinked documents create ambiguity about what should be considered attachments in legal proceedings.

5. Versions - Dynamic documents with multiple versions and changing permissions complicate historical record-keeping.

6. AI as a data source - Generative AI tools are creating new categories of data artifacts that may require legal preservation and review.

7. AI as a solution - Artificial intelligence tools offer potential solutions for managing complex emerging data challenges.

8. Continuous evolution - Rapid platform updates and new technologies make it difficult to establish consistent standards.

9. Cloud-first forensics - Traditional forensic methodologies require adaptation for cloud-based data collection and preservation.

10. Modern information governance - Existing data retention and legal hold policies need updating for current technological realities.

A Decade of Digital Transformation

The report traces the evolution from 2015, when 30% of corporate data had migrated to cloud storage, through the pandemic-driven explosion in collaboration tool adoption in 2020, to today’s complex multi-platform environment.

Key milestones include the 2017 launch of Microsoft Teams, significant court rulings on hyperlinked content in 2021 and 2024, and the emerging focus on AI-generated content as a new category of legal evidence.

The research suggests organizations need to fundamentally rethink their approach to data governance and legal readiness. This includes updating data mapping, legal hold procedures, and retention policies to account for cloud environments, chat platforms, and AI-generated content.

“Organizations must be aware of the new realities of emerging data sources in the context of legal and compliance, and develop an understanding of how to avoid common pitfalls that may quickly derail a case,” the report concludes.

The full report is available here

Two Decades. One Vision.

Smarter Information Management

ELO Digital Office

AT THE ♥ OF YOUR BUSINESS

For 20 years, ELO has empowered Australian businesses to control their information. Our enterprise content management platform is purpose-built for compliance, scalability, and growth.

20 YEARS

With the Privacy and Other Legislation Amendment Act 2024 (Cth) and further upcoming amendments - the stakes for compliance are rising.

Add sector-specific obligations for infrastructure and cyber-risk governance and a robust ECM platform like ELO is not just a nice-to-have - it's a necessity.

ELO scales effortlessly across cloud, on-premises, and hybrid deployments. Our platform combines artificial intelligence and low-code technology to connect information, automate processes, and empower your people across locations and divisions.

ELO integrates with your existing IT landscape, Microsoft 365, and third-party systems - bridging content from multiple sources into one governed platform.

Trusted by thousands worldwide and serving Australian businesses for two decades. Made in Germany, built for the world - ELO delivers enterprise-grade ECM with localised support and global innovation.

Celebrating 20 Years in Australia, and we're just getting started.

elo.com/en-au/ | info@elodigital.com.au | 1300 066 134

What my mango tree taught me about AI

BY Nicholas Fripp

It's hard to attend a meeting today without AI taking centre stage. Everybody is talking about what's next for AI, what benefits it will provide, which tools it is going to be available within, and how it is going to make your life easier, or even worse, take your job!!

In the corporate world, AI is now becoming so widely used to help generate ideas, develop draft documents, provide strategies and suggested next steps. Many software suppliers are even including it within their software to help users gain deeper insights from the data stored within these solutions, i.e., "provide me a summary of how we have responded to this type of customer question in the past", or "what is our latest policy on bereavement leave".

The technology has become very advanced, and it is almost like having your own personal secretary to help you with your daily tasks.

But… what happens when you have large amounts of records and data in your systems? Are you really getting the accurate results that you thought you were? Is AI working with you or against you?

I live in sunny Queensland, and in my backyard is an extremely large and almost oversized mango tree. This summer was a crazy mango season for our tree, and we had more mangoes that we could possibly handle.

We ate as many of them as we could. We gave away more kilograms of them then I can count. It reached a point where we had to get creative to ensure no mango was wasted.

With a little help from Google, we made sorbet, chutneys, jams, cakes… if you can think of a way to cook or bake a mango, I guarantee we tried it.

It made me think how much this tree relates to what we hold in our corporate systems and what we are trying to achieve with the "magic" of AI without any real structure, plans or proper management.

There was so much data (fruit in this case) - if I had used our standard AI approach and uploaded a photo to AI and asked for help, it would have provided me some good results (recipes, ideas etc.), but it would have also provided some terrible results too.

At a glance, you can see all the heathy mangoes on the tree, but the data would also show you that there were also many mangoes on the ground. However, it would not be smart enough or have the human insight to know that they are not usable, they are bruised, smashed or half eaten by the possums.

It would tell me that I had excess mangoes on the ground and what I could make from them, it may also have discovered the dog ball in the mix and thought it was an orange and start to give me information about use of oranges as well.

AI without governance: a fast track to risk

Many organisations have made in AI available without much thought, because it was touted as the future and something everyone should embrace. However, many of these organisations also saw the instantaneous downside of AI.

Within weeks of trials, they discovered that AI could find EVERYTHING that they didn't realise was even still there, or that didn't have appropriate security permissions.

Adding AI or even enterprise searching to your Microsoft worlds, Google spaces, and your records and information management systems is a great idea, and a great way to improve decision making; however simply adding them on top of a system with no pre-planning will provide your organisation with unfavourable results.

As records and information management specialists, I highly encourage you to utilise these technologies but be prepared to take a seat at the table to define how they are used, what your data looks like, how to protect the data and how to get the most out of the AI / searching tools.

My key tips

Here are some of my key tips as records and information professionals of how to prepare your systems for this technology.

Retention and Disposal - Use your retention and disposal tools and correctly - dispose of records and information when they are due for disposition. The old saying used to be, storage is cheap and we can always buy more storage. However, having all that information within your AI's reach will provide you with a lot of out of date, incorrect or even harmful results. There have been many cases in the press about AI agents providing back results to the public or to staff on out-of-date procedures, policies etc. The last thing you want is the AI engine telling someone that the procedure from 1991 contains the most useful information on seatbelts… In a recent conversation about AI and removing old information, there was a comment of 'You have to clean your house before you can have a party', which can easily be translated for how we prepare our corporate data for the AI and enterprise search party!

Reduce your data pool – Just because you have 40 million records stored within your repositories, that doesn't mean you will get better results out of AI. In fact, the more data you point the AI at, the more skewed your results will become, and the longer it will take to provide you the correct information. Often, you only require results relevant to the last 2-3 years of information or pertaining to particular topics. By reducing the data, you allow the AI or search engine to index smaller parameters that will provide you with much more accurate results.

Storing high value records in context – Appropriately storing records within structures and in context is one of the most effective way to reduce AI hallucinations, whilst also improving your overall records management. Applying additional metadata against records provides greater context to the AI solution so that it can better understand the meaning / purpose of the information instead of just words on a page. I once heard someone say "Context turns data into information, where that information can then be turned into knowledge"

Security frameworks – Open security is a very powerful and popular security framework and is certainly better than locking everything down to yourself so that no one can ever find it. However, when you are opening / allowing access to AI or other search engines with much deeper reach than human searching, you need to ensure that there are appropriate security measures in place to protect records that require privileged access. Perhaps the better principle is 'open by default but secured if required'.

Be very specific in your requests – While you would assume from the title of the solution, i.e. Artificial Intelligence, it sounds like is should know everything, in reality it is very much like a toddler that needs to be provided with step-by-step details and instructions. There was a video going around a few years back where the father asked their child to write instructions on how to make peanut butter toast, the father purposely took the instructions as literal as he could to prove a point, i.e., if the instructions said place the peanut butter on the bread, he quite literally placed the peanut butter jar on the bread and served it back to his child, and repeated similar activities until the child learnt to be very specific in how to make the sandwich. Much like this, AI works best where you provide detailed prompts and very granular instructions on how it should search, what persona it should be searching under, what are the parameters, and in what format you would like the results.

If you can take the time to properly define and setup the rules of how your organisation utilises AI, and you work to ensure your data is appropriately managed, then both you and your organisation will be able to enable AI to work for you and not against you.

With over 20 years of experience in the documents and records management industry, iCognition's Nicholas Fripp brings deep expertise in eDRMS solutions — including Micro Focus Content Manager — along with a strong track record spanning customer-side and consulting roles across diverse industries.

Global Framework Secures AI Systems

The European Telecommunications Standards Institute (ETSI) has published the first globally applicable cybersecurity standard for artificial intelligence systems, establishing baseline security requirements across the full AI lifecycle.

ETSI EN 304 223, published on 15 January 2026, provides a structured framework to protect AI models and systems from sophisticated cyber threats including data poisoning, model obfuscation and indirect prompt injection. The standard has been formally approved by National Standards Organisations voting, giving it broader international scope and authority across global markets.

The European Standard defines 13 principles and requirements across five phases - secure design, secure development, secure deployment, secure maintenance and secure end of life. Each phase aligns with internationally recognised AI lifecycle models, ensuring consistency and interoperability with existing standards and guidance.

The standard provides a practical baseline for securing AI systems throughout their lifecycle. Organisations implementing AI-driven workflows for compliance processes, records management or business automation can use the standard to ensure security is embedded by design.

ETSI standards provide an internationally recognised framework that aligns with the standards-led approaches adopted by both Australia and New Zealand for AI governance.

Australia participates actively in ETSI through several government bodies including the Australian Signals Directorate, the Department of Home Affairs and CSIRO, who are ETSI members. Australia frequently adopts or adapts ETSI standards for local use, with ETSI standards forming the backbone of several critical cybersecurity regulations.

The most significant recent development is the Cyber Security Act 2024, which mandates specific security requirements for smart devices. The Cyber Security (Security Standards for Smart Devices) Rules 2025 take effect on 4 March 2026, requiring manufacturers and suppliers of internet-connected consumer devices to meet mandatory security standards based on ETSI EN 303 645.

"ETSI EN 304 223 represents an important step forward in establishing a common, rigorous foundation for securing AI systems," said Scott Cadzow, Chair of ETSI's Technical Committee for Securing Artificial Intelligence.

"At a time when AI is being increasingly integrated into critical services and infrastructure, the availability of clear, practical guidance that reflects both the complexity of these technologies and the realities of deployment cannot be underestimated."

The standard builds on ETSI's earlier Technical Specification TS 104 223 but carries greater weight as a European Standard. It acknowledges that AI represents a distinct cybersecurity challenge compared to traditional software, requiring cyber defences that account for unique AI characteristics and vulnerabilities.

ETSI EN 304 223 covers AI systems incorporating deep neural networks, including generative AI, and is developed for systems intended for real-world deployments. The standard provides stakeholders throughout the AI supply chain - from vendors to integrators and operators - with a clear baseline for AI security.

An upcoming Technical Report, ETSI TR 104 159, will apply the ETSI EN 304 223 principles specifically to generative AI, focusing on deepfakes, misinformation, disinformation, confidentiality risks and copyright concerns.

SharePoint Compliance Tools April Retirement

Microsoft will retire support for several legacy SharePoint Online information management features starting April 2026, forcing organisations to migrate to modern Microsoft Purview solutions.

The retirement affects Information Management Policies, In-Place Records Management and deletion-only document policies. While the features will continue to function, Microsoft will no longer provide support after the April deadline.

Organisations using these legacy features must migrate to Microsoft Purview Data Lifecycle Management and Purview Records Management before support ends. Microsoft will not automatically migrate existing configurations.

"Microsoft won't automatically migrate your older information management and records management features in SharePoint for Microsoft 365," Microsoft stated in the announcement. "If you choose not to migrate to supported features, the older features might no longer be supported."

The retirement follows Microsoft's long-term plan to streamline compliance and data lifecycle management. The company first announced the deprecation timeline in November 2023.

Records managers and compliance teams now have until April 2026 to plan and execute their migration strategies. Microsoft recommends organisations review current use of legacy features and confirm compliance goals.

The migration preparation includes auditing retention schedules for duplicates or outdated policies. Organisations must also confirm applicable licences for Purview features and verify functionality using tools like Policy Lookup.

"Between now and the deprecation date of the feature, you have the flexibility to migrate your scenarios on your own schedule," Microsoft noted in the announcement.

AI Chatbots Must Earn Citizen Trust

Government agencies pursuing conversational AI for service delivery face significant citizen adoption barriers, with only 2% of citizens choosing AI chatbots as their first-choice channel, according to new Gartner research.

The report reveals 46% of government CIOs believe conversational AI (CAI) will play a critical or significant role in future digital service delivery. However, citizen hesitancy threatens this digital transformation vision.

"Citizens today don't consider CAI channels as their first choice for government interactions, but this is likely to change as they become more common across all industries," said Dean Lacheca, VP Analyst at Gartner and author of the report.

Among citizens who have used government conversational assistants, 25% are not willing or only slightly willing to try the channel again. The majority (57%) have never used such services.

Gartner identifies key risks deterring government adoption: AI hallucinations generating incorrect information, inadvertent exposure of sensitive citizen data, workflow disruptions, and reputational damage from unprofessional responses.

"A citizen's first user experience with a government CAI will significantly influence their future willingness to use the channel, so it is essential to design it with a human-centred approach from the outset," Lacheca said.

Evolution Beyond Traditional Chatbots

For over a decade, governments have implemented chatbots as part of multichannel service delivery strategies. These systems, designed for straightforward inquiries and automated responses, have delivered mixed results.

Traditional chatbots often suffer from negative citizen perceptions due to outdated technologies or poor implementation. They frequently require extensive human intervention to manage escalated issues.

Recent generative AI advancements have enabled CAI technologies to shift from simple query handling to sophisticated interactions that mimic human conversation patterns, propose next-best actions and solve problems.

The convergence of advanced CAI technologies and shifting citizen expectations is changing the benchmark for government service delivery.

Risk Mitigation Strategies

Government CIOs must implement robust risk mitigation strategies including active monitoring, clear escalation paths to human agents, and rule-based AI monitoring alongside generative AI components, according to Gartner.

"Poorly thought-through implementations will undermine trust and fail to meet citizen expectations for quick, low-effort service experiences that are available 24/7 and are personalised," Lacheca said.

The research recommends pilot programmes with internal users before citizen-facing deployment, co-creation with citizen groups, and continuous content curation of AI-optimised knowledge repositories.

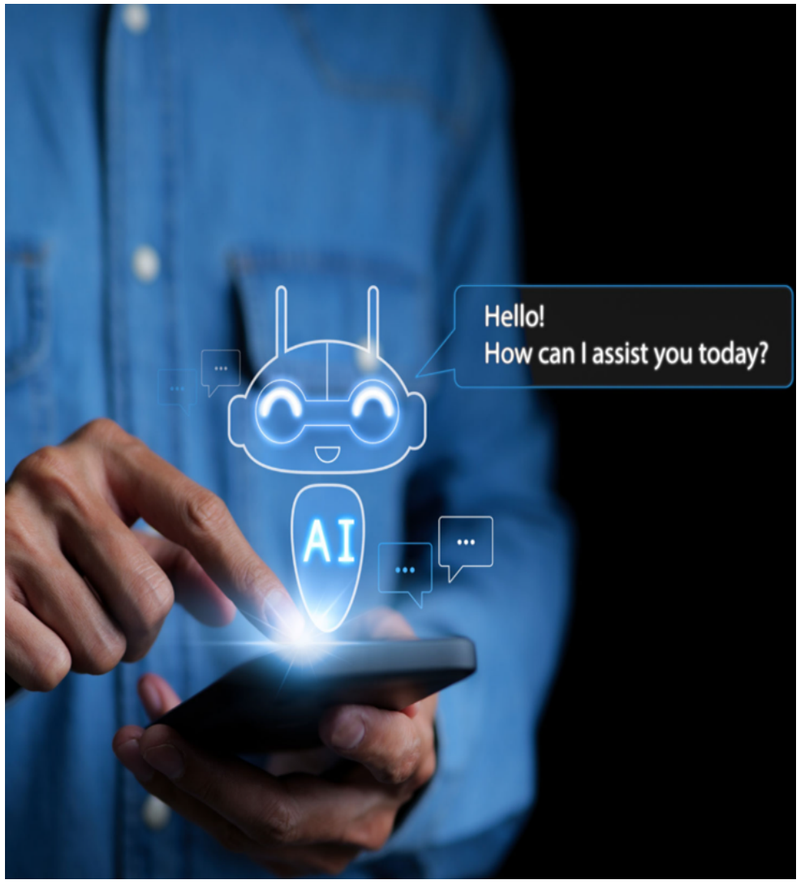

Survey Methodology

The findings draw from two Gartner surveys conducted in 2025. The Transition to Digital Government Survey, conducted July through September, surveyed 138 government respondents from North America, EMEA, and Asia/Pacific screened for involvement in digital solution adoption.

The Citizen Experience Survey investigated citizen preferences among respondents aged 18 or older who interacted with government services in the past 12 months. Respondents resided in the United States, Canada, United Kingdom, Netherlands, Germany, Sweden, Norway, France, Spain, or Australia.

Dean Lacheca is VP Analyst in Gartner's CIO business and technology insights group. Prior to joining Gartner, Mr Lacheca worked in the public sector space for more than 20 years including 11 years at the government owned worker's compensation insurer in Queensland.

Budget Cuts Hit Privacy Teams Hard

Privacy team sizes have plummeted by more than one-third globally, with the median dropping from eight staff to five, according to ISACA's State of Privacy 2026 report.

The survey of more than 1,800 privacy professionals conducted in September 2025 reveals mounting pressure on organisations struggling to maintain compliance amid resource constraints and rapid technological change.

Half of respondents anticipate privacy budget decreases in the next 12 months, while fewer than one-quarter expect increases. Technical privacy roles remain more understaffed than legal and compliance positions, with 47% reporting understaffing compared with 37% for legal and compliance roles.

"Privacy teams are shrinking," the report states. "The median privacy staff size of survey respondents is five, down from eight last year."

Jamie Norton, Vice Chair of the ISACA Board, said privacy teams across Oceania are being stretched at a time when expectations continue to rise. "Many organisations are asking small privacy teams to manage complex compliance obligations, emerging technologies like AI, and growing breach risk all at once," said Mr Norton.

"Lower budgets can mean that organisations risk falling behind regulatory expectations as scrutiny continues to intensify. When investment doesn't keep pace, privacy risk quickly becomes a broader business and governance issue."

Confidence in meeting compliance requirements has also declined, with fewer than half of respondents (46%) reporting they are very or completely confident in their privacy team's ability to achieve compliance with new privacy laws and regulations.

The findings come as organisations grapple with an increasingly complex regulatory landscape, including new privacy laws across multiple jurisdictions and emerging requirements around artificial intelligence governance.

Survey respondents identified lack of training or poor training as the most common privacy failure at 51%, up from 47% last year. Not practising privacy by design ranked second at 50%, representing a nine percentage-point increase from 2025.

Technical privacy professionals face particular challenges, with 54% of respondents citing technical expertise as the biggest skill gap, followed by experience with different types of technologies and applications at 52%.

The rapid evolution of technology emerged as the top stressor for privacy professionals at 71%, up eight percentage points from last year. Compliance challenges (62%), resource shortages (61%) and competing priorities (56%) rounded out the leading stressors.

Despite these pressures, organisations with strong board support and strategic alignment fare better. The report found that 29% of respondents whose boards adequately prioritised privacy anticipated budget increases, compared with just 14% whose boards did not prioritise privacy.

"A CPO is an important role that can advocate for privacy teams and initiatives," the report notes. Respondents whose organisations had a chief privacy officer were more likely to feel their board adequately prioritised privacy and more confident in ensuring data compliance.

The survey also explored artificial intelligence adoption, with 13% of respondents currently using AI for privacy-related tasks and 38% planning to adopt AI within 12 months. However, AI adoption correlated strongly with organisational maturity, with 41% of current AI users reporting they always practice privacy by design.

Organisations with privacy strategies aligned to broader business objectives demonstrated better outcomes across multiple metrics, including less understaffing, more optimistic budget forecasts and higher rates of practising privacy by design.

ISACA surveyed 47,600 constituents holding the Certified Data Privacy Solutions Engineer (CDPSE) designation, Certified Information Security Manager (CISM) designation or having "privacy" or "data protection" in job titles.

The full State of Privacy 2026 report is available HERE.

State of Privacy 2026

AI Safety Takes Backseat in National Plan

Australia's National AI Plan has drawn sharp criticism from legal and academic experts who warn the government's decision to rely on existing legislation leaves organisations exposed to emerging risks in high-stakes automation and automated decision-making systems.

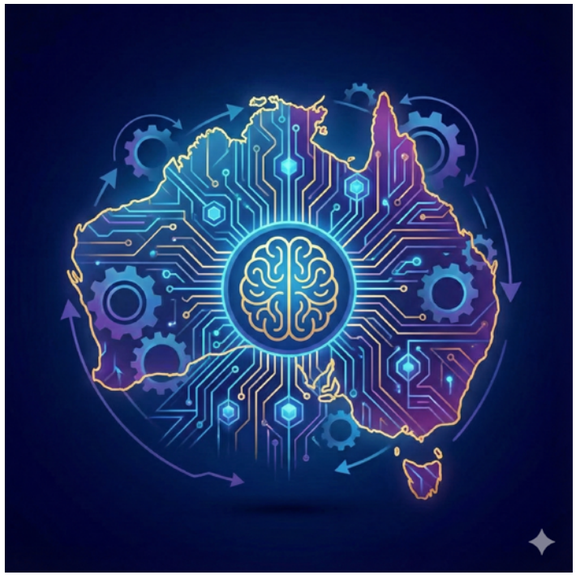

The plan abandons previously proposed mandatory guardrails for high-risk AI systems, instead establishing a light-touch regulatory framework built on technology-neutral laws covering privacy, consumer protection and workplace safety.

This approach has divided experts, with critics arguing existing frameworks are inadequate for managing AI-specific governance challenges while supporters praise the focus on innovation and economic opportunity.

"The absence of new AI-specific legislation means Australia still needs clearer guardrails to manage high-risk AI," said Associate Professor Sophia Duan from La Trobe University. "Trustworthy AI requires more than voluntary guidance."

Dr Rebecca Johnson, an AI ethicist at the University of Sydney, said relying on technology-neutral laws fails to address the unique autonomy of modern AI agents.

"It's like trying to regulate drones with road rules: some parts apply, but most of the risks fly straight past," Johnson said. "AI agents don't just generate text; they carry out tasks… That is a fundamentally different safety landscape"

However, Professor Niloufer Selvadurai from Macquarie Law School praised the strategy for avoiding the complications of creating new laws for evolving tech.

"Given the complex and diverse applications of AI, I think this nuanced approach, premised on a regulatory gap-analysis, is to be welcomed," Selvadurai said.

The plan focuses on three objectives - capturing economic opportunities through data centre investment, spreading benefits via workforce training, and keeping Australians safe through enhanced oversight. However, the government will rely on existing legal frameworks rather than introduce standalone AI legislation, a reversal from earlier consultation documents that proposed mandatory controls.

For organisations managing compliance processes and digital transformation, the decision creates uncertainty around governance requirements for AI systems used in automated decision-making, document processing and data analytics. The plan provides no specific guidance on how existing privacy, discrimination and consumer protection laws will apply to AI deployments in records management, information governance or business intelligence applications.

Dr Malcolm Thatcher, a digital risk governance specialist, said the lack of an AI-specific regulatory framework leaves organisations exposed.

"The Australian Government has dropped the ball on providing an AI-specific regulatory framework for the safe use of AI," he wrote. "This leaves organisations who don't fully understand AI and all its complexities including potential interaction with other laws, exposed to this increasing AI risk."

Legal experts from Bird & Bird noted the plan signals heightened regulatory scrutiny without creating new legal obligations. "For organisations operating in or into Australia, this Plan sets the direction of travel for investment, regulation, workforce policy and government procurement over the rest of this decade," the firm stated.

"Expect more public investment and procurement activity, alongside heightened expectations for responsible governance and transparency."

The plan establishes an AI Safety Institute with $A30 million in funding to monitor AI harms and advise on regulatory interventions. However, the institute will operate in an advisory capacity without statutory powers, relying on existing regulators to enforce technology-neutral legislation.

King & Wood Mallesons characterised the approach as reducing immediate compliance burden while requiring organisations to establish responsible AI governance frameworks.

"Although this does reduce the compliance burden on companies, they will still need to establish responsible AI Governance," the firm advised.

The regulatory gap analysis extends to critical areas for information managers and compliance professionals. Privacy law reforms remain incomplete, with the second tranche of amendments to the Privacy Act still without clear timeframes.

Copyright protections for AI training data lack resolution, following the government's October 2025 decision not to introduce a broad text-and-data-mining exception.

For organisations deploying AI in information management, records processing or compliance workflows, the plan emphasises transparency and documentation requirements.

The National AI Centre's "AI6" governance practices - released in October 2025 - provide baseline expectations for risk assessment, human oversight and accountability structures.

Strong Risk Frameworks Key To GenAI Success

Organizations must elevate risk management capabilities to achieve meaningful returns on generative AI investments, according to new research from Parker & Lawrence examining AI adoption in risk and compliance functions.

The report identifies an "ROI paradox" where 95% of organizations see zero return despite $US30-40 billion in enterprise GenAI investment.

Firms default to low-risk use cases like summarization and chatbots that deliver limited strategic value while carrying GenAI's relatively high costs.

"Most organizations fall on the wrong side of the GenAI Divide. Adoption is high, but disruption is low," MIT researchers noted in findings cited by the report.

Survey data from 224 senior risk and compliance professionals across the UK and US revealed 98.2% report GenAI presents new or increased challenges.

Financial crime, compliance management and cybersecurity face the greatest financial impact from both GenAI opportunities and risks.

Less than half of respondents feel "very confident" in their organization's ability to control GenAI risks. Data quality and availability issues topped implementation barriers at 45.3%, followed by cost concerns.

The research found organizations experiencing highest ROI are more likely to require technical staff have GenAI risk mitigation skills (68% versus 34% for low-ROI organizations) and maintain clear processes for embedding risk controls (44% versus 23%).

"Getting better at being able to manage AI risks allows you to deploy those solutions quicker and with a lot more confidence," said Kristof Horompoly, VP of AI Risk Management at ValidMind.

Parker & Lawrence evaluated 620 GenAI applications across risk and compliance workflows, identifying 47 transformative use cases primarily involving advanced reasoning capabilities.

Data generation and structuring showed greatest breadth of applications, while constraint-based evaluation and hypothesis generation delivered most transformative impacts.

The report examined seven technology vendors demonstrating effective GenAI integration in areas including fraud detection, transaction monitoring, compliance assurance, regulatory obligations mapping, identity management, operational resilience and trade surveillance.

Industry experts interviewed emphasized that effective governance must move beyond principles to embedded practice.

"AI governance cannot be fulfilled by simply publishing a white paper. It must be embedded across controls, practices, and processes," said Anna Nicolis, risk management consultant at Shapes First.

The research identifies three critical priorities: applying GenAI to decision-critical problems rather than incremental efficiencies, adopting holistic approaches addressing multiple concurrent risks, and treating risk management as strategic infrastructure rather than bottleneck.

Organizations at the "risk optimization" stage - representing approximately 22% of firms - actively scale AI across core functions while building frameworks enabling higher-value use cases through formalized governance and standardized validation protocols.

Name and Shame Seeks AI Assistant

The NSW Food Authority has issued a tender for an artificial intelligence tool to automate the assessment of penalty notices, aiming to replace a manual process currently requiring a full-time employee earning $115,000 annually.

It is looking for an AI-powered solution to verify the accuracy of penalty notices before publication on the Food Authority's Name and Shame website, with submissions closing 9 February 2026.

The initiative addresses significant compliance and operational challenges within NSW's food safety enforcement system. In the 2024/25 financial year, the Food Authority assessed 1,178 penalty notices issued under the NSW Food Act 2003, with an estimated 40 per cent containing errors or incomplete information requiring manual verification and follow-up with issuing agencies.

The current manual process involves opening each penalty notice record in the Byte regulatory business system, checking scanned documents for accuracy, verifying legal entity details match offence codes and fine amounts from the Fixed Penalty Handbook, and conducting extensive back-and-forth correspondence with 128 NSW councils and enforcement agencies to resolve discrepancies or obtain missing information.

Penalty notices arrive in both handwritten and electronic formats, requiring assessors to verify correct legal entity identification, appropriate offence codes and fine amounts, breach clause citations, complete address information, and premises type classifications before determining publication eligibility through a decision matrix based on priority classification and breach severity.

The Food Authority's Name and Shame website recorded 770,000 hits in 2024/25, creating heightened brand and reputation risks if incorrect information is published. Section 133A-H of the Food Act 2003 mandates the Food Authority publish penalty notice information in good faith and remove improperly served notices from the register.

The proposed AI solution will use image recognition technology to categorise penalty notice data before ingestion into the existing Byte workflow engine, complementing or replacing manual assessment steps where possible. The stand-alone tool will integrate with Byte, which manages food licensing, biosecurity authorisations, inspections, enforcement actions and the Name and Shame register.

Revenue NSW currently provides penalty notice files to the Food Authority's SFTP site weekly, triggering automated imports into Byte. The AI tool aims to supplement this workflow by automatically identifying incorrectly issued notices that cannot be published, flagging notices with missing verification information, and validating complete and correct notices ready for publication.

The successful supplier must train internal specialists to provide ongoing technical support in-house, ensuring knowledge transfer and reducing dependence on external vendors for system maintenance.

Project objectives include ensuring watertight and consistent assessment procedures, providing due diligence assurance for published information, significantly reducing staff time and resource burden, minimising human error risk, delivering cost savings, and creating equitable processes for all offenders.

The project timeline spans 20 February to 19 June 2026, including supplier engagement, scoping workshops, system design and development, testing and pilot implementation.

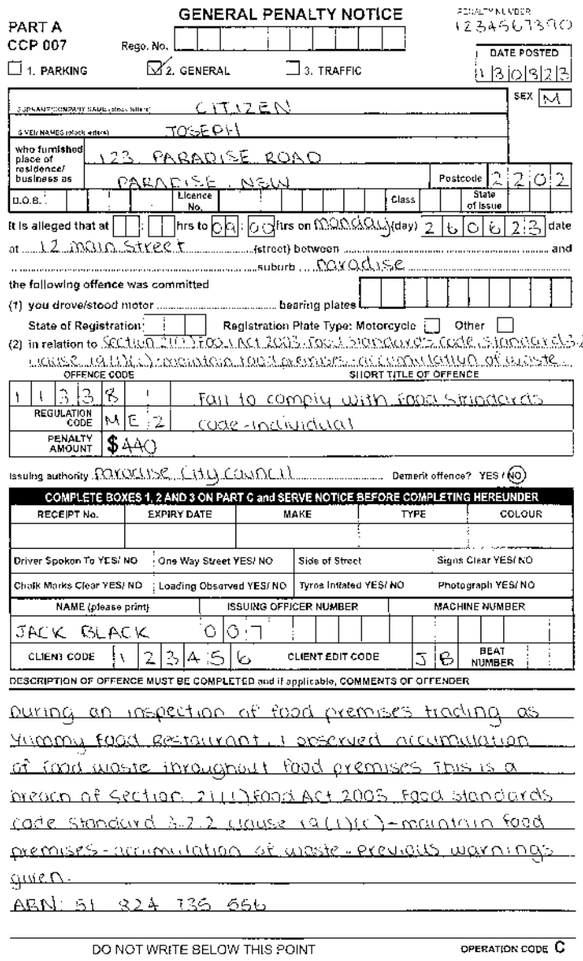

GENERAL PENALTY NOTICE

PART A
CCP 007

Rego. No.

PENALTY NUMBER 1234567890

☐ 1. PARKING ☑ 2. GENERAL ☐ 3. TRAFFIC

DATE POSTED 13/02/23

SURNAME/COMPANY NAME (block letters) CITIZEN SEX M

GIVEN NAMES (block letters) JOSEPH

who furnished place of residence/ business as 123 PARADISE ROAD PARADISE NSW Postcode 2202

D.O.B. Licence No. Class State of Issue

It is alleged that at 09:00 hrs on Monday (day) 26/06/23 date

at 12 Main Street (street) between and

suburb Paradise

the following offence was committed

(1) you drove/stood motor bearing plates

State of Registration: Registration Plate Type: Motorcycle ☐ Other ☐

(2) in relation to Section 21(1) Food Act 2003 Food Standards Code Standard 3.2.2 clause 19(1)(c) - maintain food premises - accumulation of waste

OFFENCE CODE	SHORT TITLE OF OFFENCE
113381	Fail to comply with food Standards
REGULATION CODE ME2	code - individual
PENALTY AMOUNT $440	

Issuing authority Paradise City Council Demerit offence? YES / NO

COMPLETE BOXES 1, 2 AND 3 ON PART C and SERVE NOTICE BEFORE COMPLETING HEREUNDER

RECEIPT No.	EXPIRY DATE	MAKE	TYPE	COLOUR

Driver Spoken To YES/ NO	One Way Street YES/ NO	Side of Street	Signs Clear YES/ NO
Chalk Marks Clear YES/ NO	Loading Observed YES/ NO	Tyres Initated YES/ NO	Photograph YES/ NO

NAME (please print)	ISSUING OFFICER NUMBER	MACHINE NUMBER
JACK BLACK	007	

CLIENT CODE 123456 CLIENT EDIT CODE JB BEAT NUMBER

DESCRIPTION OF OFFENCE MUST BE COMPLETED and if applicable, COMMENTS OF OFFENDER

During an inspection of food premises trading as Yummy Food Restaurant, I observed accumulation of food waste throughout food premises. This is a breach of Section 21(1) Food Act 2003, Food Standards Code Standard 3.2.2 clause 19(1)(c) - maintain food premises - accumulation of waste - previous warnings given.

ABN: 51 824 735 556

DO NOT WRITE BELOW THIS POINT OPERATION CODE C

A handwritten penalty notice

Patient Data at Risk in NSW Hospitals

An audit of 4 out of 15 NSW Local Health Districts (LHDs) by the state's Auditor-General found they failed to meet minimum cyber security requirements, leaving clinical systems vulnerable to attacks that could disrupt healthcare delivery.

The report, completed in July 2025 but withheld until 19 December to allow remedial action, found only one Local Health District has an incomplete cyber security plan. The remaining audited Local Health Districts do not have a cyber security plan at all.

"NSW Health is not effectively managing cyber security risks to clinical systems that support healthcare delivery in Local Health Districts," the report states. It does not name the four LHDs that were audited.

Auditor-General Bola Oyetunji presented the report confidentially to Parliament on 9 July 2025. "I determined that it was not in the public interest to make the report public at that time," Oyetunji said in an addendum. The five-month delay allowed NSW Health to establish a taskforce and progress responses to recommendations before public disclosure.

"Systemic non-compliance with NSW Government cyber security requirements, including maintaining adequate cyber security response plans, business continuity planning and disaster recovery for cyber security incidents, means that Local Health Districts could not demonstrate that they are prepared for, or resilient to, cyber threats," the report warns.

"Local Health Districts are not adequately prepared to respond effectively to cyber security incidents. This exposes the risk that a preventable cyber security incident could disrupt access to healthcare services."

Clinical staff routinely disregard cyber security controls, creating what the report describes as "normalisation of non-compliance" driven by tensions between clinical urgency and security protocols.

The audit observed multiple violations across all audited districts, including storing patient information outside secure systems and leaving computers logged in when unattended.

"Despite being aware of clinical staff's systemic non-compliance with cyber security controls, the audited Local Health Districts have not undertaken work to assess the effectiveness of the controls," the report states.

"eHealth NSW has not clearly defined and communicated its roles and the expected roles of Local Health Districts for cyber security," according to the report.

The report found Local Health Districts spent an average $A421,000 on cyber security in 2023-24, representing just two per cent of ICT expenses against a benchmark of nine per cent.

None of the audited districts demonstrated consideration of cyber controls beyond minimum requirements, despite handling large volumes of sensitive personal and health information.

"NSW Health reported that healthcare is the most breached industry in Australia, and health data is 50 times more valuable than credit card data on the dark web," the report notes.

Business continuity and disaster recovery plans do not adequately address cyber security incidents, leaving districts unprepared for attacks that could compromise patient care.

The report found eHealth NSW conducted its first-ever test of the overall NSW Health cyber security incident response plan in October 2024, but it excluded clinical involvement.

"Clinicians did not participate in the exercise and were not consulted on the findings or recommendations resulting from the exercise," the report states.

NSW Health's 2024 attestation to Cyber Security NSW aggregated compliance across 32 organisations, obscuring risks within individual Local Health Districts.

"The attestation obscures the cyber security risks that exist for each Local Health District," the report warns. "NSW Health and Cyber Security NSW may not fully understand cyber security risk in this part of the health system."

The audit made six recommendations for implementation by December 2025, including collating and validating compliance information, finalising cyber security roles and developing clear guidance for balancing clinical service delivery with security requirements.

NSW Health Secretary Susan Pearce accepted all recommendations. "NSW Health is committed to ensuring that our system is safe from cyber security threats and that the sensitive information we hold is safeguarded," Pearce said in her response.

A dedicated Cyber Security Uplift Program has been established to enhance cyber resilience and ensure compliance with the NSW Cyber Security Policy and Security of Critical Infrastructure Act 2018.

After the report is presented to the NSW Parliament, it is usual for the entity's Audit and Risk Committee/Audit Risk and Improvement Committee to monitor progress in implementing recommendations.

In addition, it is the practice of NSW Parliament's Public Accounts Committee to conduct reviews or hold inquiries into matters raised in performance audit reports. The reviews and inquiries are usually held 12 months after the report is received by the NSW Parliament. These reports are available on the NSW Parliament website.

The report is available HERE.

Agencies Fail OAIC Transparency Test on Decision Systems

Only 17 per cent of Australian Government agencies disclose their use of automated decision-making (ADM) systems despite legal requirements to inform the public, according to a new report from the Office of the Australian Information Commissioner (OAIC).

The review of 23 agencies authorised to use ADM found 74 per cent could not be identified as using the technology through publicly available information. No agency has published guidelines or policies on ADM use.

"Proactively publishing clear information about automated decision-making is essential to building trust and ensuring accountability," Australian Information Commissioner Elizabeth Tydd said.

"Through this Report we can encourage agencies to deliver greater community awareness and satisfaction about how government decisions are made."

The OAIC's report, Automated Decision-Making And Public Reporting Under The Freedom Of Information Act, follows a desktop review conducted in October 2025 of agency websites, AI transparency statements and Information Publication Scheme (IPS) materials.

ADM refers to the use of technology, commonly called a 'computer program' in Commonwealth legislation, to automate decision-making processes. It is used across government in areas including social services, taxation, aged care and veterans' entitlements.

The review found all 23 agencies publish IPS-related information on their websites. However, only four agencies - the Australian Taxation Office, Services Australia, Department of Veterans' Affairs and Department of Health, Disability and Aging - disclosed that they use ADM in decision-making processes.

A further nine agencies made reference to ADM in their IPS information but did not confirm whether they used it. Ten agencies made no mention of ADM despite having statutory authority to use it.

Commissioner Tydd said the findings highlighted opportunities for improvement in meeting Freedom of Information Act 1982 obligations.

"Information about decision-making and the exercise of agencies functions is important information for the Australian community," she said. "It improves integrity, accountability and trust."

The report identified several case studies illustrating the transparency gap. One regulatory agency authorised to use ADM has an AI transparency statement saying it does not use AI without human oversight. However, external evidence found the agency employs ADM to calculate fees in its online portal, with the explanatory statement for the legislative instrument explicitly confirming this use.

Another case study described an agency that mentioned ADM only in a data strategy report, stating it was "embracing automation and artificial intelligence, which allows it to make decisions based on data, in a more timely manner." The agency did not elaborate on how decisions are made or whether any are based solely on automated processes.

"Without published guidelines, policies or procedures," the report stated, "we did not categorise any agency as 'better practice'. This is despite 4 agencies acknowledging that they use ADM and a further 2 that had external evidence suggesting that they also used ADM."

The review follows the 2023 Robodebt Royal Commission, which found the scheme's heavy reliance on ADM for income averaging did not comply with Social Security Act provisions. The Commission recommended legislative reform to introduce consistent legal frameworks where automation in government services can operate, including clear paths for review and plain language explanations on departmental websites.

The OAIC report makes four key recommendations for agencies authorised to use ADM:

1. Publish information as part of the IPS, including the statute granting power and whether ADM is utilised to provide information and services.

2. Clearly state the types of ADM used to make automated decisions, not just AI - from simple calculators to machine learning.

3. Publish both a list of decisions where ADM is used and relevant, easy-to-understand examples.

4. Publish policies that clearly set out principles for when and how ADM is used to make decisions affecting members of the public.

Commissioner Tydd confirmed the OAIC will begin consultation to update the Information Commissioner Guidelines as a priority in 2026.

"The OAIC will update Part 13 - Information Publication Scheme of the FOI guidelines so that ADM is expressly included as an example of 'operational information'."

The report identified good practices among the four agencies that disclosed ADM use. The Department of Veterans' Affairs disclosed ADM through multiple reference points including privacy policy and data collection notices, listed services where ADM may be used, and listed all relevant legislation.

Services Australia provided a dedicated automation and AI webpage with technical information on three broad types of automation used. The ATO provided an easy-to-read case study about ADM in relation to franking credits.

The Privacy and Other Legislation Amendments Act has added requirements for entities to include in privacy policies information about whether ADM has substantially and directly made decisions that significantly affect individual rights or interests.

The full report and recommendations are available HERE.

Is this glass square the long, long future of data storage?

By Alex Fuerbach, Macquarie University

Scientists at Microsoft Research in the United States have demonstrated a system called Silica for writing and reading information in ordinary pieces of glass which can store two million books' worth of data in a thin, palm-sized square.

In a paper published in Nature, the researchers say their tests suggest the data will be readable for more than 10,000 years. The new system, called Silica, uses extremely short flashes of laser light to inscribe bits of information into a block of ordinary glass.

These pulses are called "ultrashort" for a reason. Each one lasts mere quadrillionths of a second (aka femtoseconds or 10^{-15} s).

To get your head around that: comparing ten femtoseconds to a single minute is like comparing one minute to the entire age of the universe. These incredibly short flashes can be used to generate even shorter bursts of light lasting attoseconds (a thousandth of a femtosecond or 10^{-18} s).

These attosecond bursts can be used to observe the motion of electrons inside atoms and molecules – and in 2023 the Nobel Prize in Physics was awarded for pioneering work in this area, to Ferenc Krausz (coincidentally my former PhD supervisor), Anne L'Huillier and Pierre Agostini.

Writing in glass

Femtosecond laser pulses also have a practical technological application. They can be used to make changes deep inside transparent materials such as glass. These lasers produce light of a wavelength that normally passes through glass without interaction. However, when ultrashort pulses of this light are tightly focused on a particular region, it produces an intense electric field that alters the molecular structure of the glass in the focal zone.

This means only a tiny three-dimensional volume, often less than a millionth of a metre to a side, is affected. This is called a "voxel", which can be made at precisely controlled positions in the glass. The idea of using laser-written voxels for three-dimensional data storage is not new. Eric Mazur and co-workers at Harvard University in the US investigated volumetric optical storage back in the 1990s. Their groundbreaking work demonstrated that permanent data structures could be inscribed into common glass using femtosecond lasers.

In 2014, Peter Kazansky and colleagues at the University of Southampton in the UK reported data storage in fused quartz glass with a "seemingly unlimited lifetime". This helped to to establish the idea of ultra-stable glass-based memory devices.

In 2024, Kazansky spun out a company called SPhotonix to commercialise what they describe as "5D glass nanostructuring". Their vision of a "5D memory crystal" even made its way into popular culture: a similar device appeared in the latest Mission Impossible film, The Final Reckoning, portrayed as a secure vault capable of containing a powerful but sinister AI.

The Silica project does not claim to have made a new scientific breakthrough. Instead the team presents the first comprehensive demonstration of a practical real-world technology. Their work brings together all the key elements of such a storage platform based on femtosecond lasers and glass. It includes encoding data, writing, reading, decoding and error correction. The work explores different strategies for reliability, writing speed, energy efficiency and data density, and involves systematic assessments of the data lifetime.

Silica looked at two main types of laser-written voxels. The first consists of tiny elongated void-like features created by laser-driven "micro-explosions" inside the glass. These allow an extremely high storage density of 1.59 gigabits per cubic millimetre.

The second type involves making subtle changes in the local refractive index of the glass. These can be written faster, using less energy – but each cubic millimetre of glass can hold less data. This method can write about 65.9 megabits per second, and the authors say this could be increased with more laser beams.

Finally, accelerated ageing experiments suggest that the written data, even in the case of the more sensitive phase voxels, could remain stable for more than 10,000 years. This vastly exceeds the lifetime of conventional archival storage media such as magnetic tape or hard drives.

The future

When I began my PhD in the late 1990s at the Vienna University of Technology, we were one of only a handful of laboratories worldwide that had the expertise to build lasers capable of generating femtosecond pulses. Today, after decades of technological development, ultrafast lasers with the reliability, power and repetition rates required for industrial use can be purchased off the shelf.

Dense, fast and energy-efficient archival data storage is an exciting real-world application of these lasers. As ultrafast photonics continues to mature, I have no doubt more applications will follow. Exciting times ahead.

Alex Fuerbach is Professor, Photonics Research Centre, Macquarie University. This article is republished from The Conversation under a Creative Commons license. Read the original article.

Speed, Identity, and AI: Redefining Cyber-security for Australian Organisations

BY Harry Chichadjian

Cybersecurity teams have always grappled with emerging technologies, increasingly sophisticated attacks and new regulations. However, as organisations accelerate AI and cloud adoption, digitise services, and adapt to evolving compliance regulations, security challenges continue to grow in tandem with the expansion of attack surfaces.

What's concerning is that threat actors are no longer waiting in the shadows to probe for weaknesses. They now prioritise immediate execution- leveraging AI, targeting identity, and monetising attacks quickly.

The latest Elastic Global Threat Report (GTR) 2025 reveals that execution-first attacks on Windows now account for 32% of all observed malicious activity, nearly doubling from last year and surpassing defence evasion as the top tactic. This shift signals that attackers are prioritising speed and impact over persistence and stealth.

For Australian businesses, the message is clear: traditional perimeter-focused, compliance-driven approaches are no longer enough. Cyber resilience must be strategic, proactive, and data-driven. The window to detect and contain threats is narrowing, identity has become the most heavily targeted asset in cloud environments, and generative AI is lowering the barrier for adversaries to launch more frequent and sophisticated attacks.

Execution First: Speed over Stealth

The GTR shows that attackers are running malicious code immediately upon entry, shifting away from stealthy persistence. This aligns sharply with the local landscape, where the Australian Cyber Security Centre (ACSC) reported over 94,000 cybercrime incidents in 2023–24, a 23% increase year-on-year.

As adversaries accelerate their tactics, the window for detection and response is shrinking. Legacy, compliance-driven approaches are no longer sufficient. Organisations must pivot to proactive threat hunting, memory safeguards, and advanced endpoint detection that can flag anomalous execution in real-time.

Another key finding from the GTR is how attacks in cloud environments are highly focused, with over 60% of incidents targeting initial access, persistence, and credential access. With cloud adoption in Australia among the highest globally, the risks are acute given that usage is projected to grow at around 12% annually, and nearly 90% of organisations are relying on multi-cloud strategies.

Compromised accounts and credentials are the leading incident type across critical infrastructure and government sectors. Protecting identity is now paramount. Phishing-resistant MFA, least-privilege access, and continuous monitoring of privileged activity are the most effective safeguards.

At the same time, browser-stored credentials and accidental source code leaks are emerging as key vectors for attack. Around one in eight malware samples are designed to steal browser data, which fuels the global access-broker economy and drives ransomware,

business email compromise (BEC), and extortion campaigns. In Australia, these threats are increasingly driving breaches, often via employees' personal or BYOD devices, making compromised credentials the largest category of reported incidents.

AI and Source Code: Escalating Threats for Australian Organisations

AI is lowering the barrier to entry for cybercrime, enabling more frequent and varied attacks. Findings from the GTR have revealed a 15.5% increase in 'Generic' threats, which are malicious files or programmes that cannot be categorised elsewhere. The increase may have been driven by the use of large language models (LLMs) to quickly generate malicious loaders and tools. Australian organisations are already feeling the impact, with 78% reporting significant disruption and over half unprepared to defend against AI-driven threats.

Another threat that can even potentially transform into a permanent risk to organisations is accidental source code leaks. Whether they are API keys, credentials, or sensitive data, such leaks create permanent, distributed exposure that attackers can exploit long after the initial mistake, carrying regulatory, operational, and reputational consequences for Australian businesses. Mitigation requires treating browsers and developer workflows as critical security boundaries, enforcing phishing-resistant MFA, hardening endpoints, securing credentials and session tokens, continuously monitoring development environments, and embedding automated detection and remediation directly into workflows.

The Path Forward for Australian Defenders

These trends may appear to be disparate, but they are in fact deeply interconnected. For example, threat actors can use AI-generated malware to steal browser credentials and gain access to a cloud account.

The risks for Australian organisations are immediate and tangible.

The path forward is clear: prioritise runtime visibility, strengthen the identity layer, treat browsers and developer workflows as critical assets, and embed AI-driven threat detection and behavioural analytics throughout operations. Organisations that act decisively can turn today's complex threat landscape into a strategic advantage- reducing risk, safeguarding critical assets, and staying one step ahead of adversaries.

Harry Chichadjian is Security Director, Elastic, ANZ

Data Readiness Gap Threatens AI Ambitions, IBM Study Finds

Only 26% of Chief Data Officers are confident their organisation's data can support AI-enabled revenue streams, despite 81% prioritising investments to accelerate AI capabilities.

The findings come from an IBM Institute for Business Value study of 1,700 CDOs across 27 countries and 19 industries. The research highlights a significant gap between AI ambitions and organisational readiness.

Data accessibility, completeness, integrity, accuracy and consistency remain major barriers to leveraging enterprise data for AI initiatives, according to the study.

IBM claims 81% of surveyed CDOs report their data strategy now integrates with technology roadmaps and infrastructure investments, up from 52% in 2023.

However, only 26% express confidence in using unstructured data to deliver business value. The study notes 81% of respondents bring AI to data rather than centralising it.

While 92% of CDOs say they must focus on business outcomes to succeed, only one-third strongly agree they can clearly convey how data facilitates business results. Just 29% have clear measures to determine the value of data-driven business outcomes, according to the study.

The research indicates deploying data for competitive advantage has become the top CDO priority, surpassing governance and security as core responsibilities.

IBM reports 84% of surveyed CDOs say their unique data products have provided significant competitive advantages, with 78% citing proprietary data as a top strategic objective.

AI agents present governance challenges

The study found 80% of leaders have started developing diverse datasets to train AI agents. However, 79% admit being early in defining how to scale and govern them.

Despite governance uncertainties, 83% believe potential benefits of deploying AI agents outweigh risks. Some 77% are comfortable with their organisation relying on AI agent outcomes.

Attracting, developing and retaining talent with advanced data skills is now a top challenge for 47% of CDOs, up from 32% in 2023.

The research indicates 77% of leaders struggle to fill key data roles. Only 53% say recruitment efforts deliver needed skills and experience, down from 75% in 2024.While 82% of CDOs say data is wasted without employee access, and 80% cite data democratisation as enabling faster organisational movement, fostering a data-driven culture remains a top strategic challenge.

The study was conducted between July and September 2025 in cooperation with Oxford Economics.

AI-Generated Data Drives Governance Framework Shift

Half of all organisations will implement zero-trust data governance by 2028 as AI-generated content floods information systems, according to Gartner.

"Organisations can no longer implicitly trust data or assume it was human generated," said Wan Fui Chan, Managing VP at Gartner. The company positions zero-trust governance as establishing authentication and verification measures to protect business outcomes.

The driver behind this predicted shift is the proliferation of AI-generated content entering training datasets for large language models. According to Gartner's 2026 CIO and Technology Executive Survey, 84% of respondents expect increased GenAI funding this year, though survey methodology and sample size were not disclosed.

Gartner warns of "model collapse" - where AI systems trained on outputs from previous models may lose accuracy.

Chan noted that regulatory requirements for verifying "AI-free" data are expected to intensify in certain regions. However, he did not identify specific jurisdictions, proposed regulations, or enforcement mechanisms currently in development.

"In this evolving regulatory environment, all organisations will need the ability to identify and tag AI-generated data," Chan said. He identified information and knowledge management skills plus metadata management solutions as essential for data cataloguing.

The prediction suggests organisations will need to implement active metadata management practices enabling realtime alerts when data requires recertification or becomes stale.

Gartner recommends organisations appoint dedicated AI governance leaders responsible for zero-trust policies, AI risk management, and compliance operations. The firm also advises forming cross-functional teams spanning cybersecurity and data analytics to conduct data risk assessments.

The company suggests building on existing data and analytics governance frameworks rather than creating entirely new systems. This includes updating security, metadata management, and ethics policies to address AI-generated data risks.

The 50% adoption prediction by 2028 represents a notably aggressive timeline.

Explore frameworks and best practices for data, analytics and AI governance at the 2026 Gartner Data & Analytics Summit, June 16 – 17 2026. Learn more HERE.

EzeScan is one of Australia's most popular production capture applications and software of choice for many Records and Information Managers. This award winning technology has been developed by Outback Imaging, an Australian Research and Development company operating since 2002. Solutions range from centralised records capture, highly automated forms and invoice processing to decentralised enterprise digitisation platforms which uniquely align business processes with digitisation standards, compliance and governance requirements. With advanced indexing functionality and native integration with many ECM/EDRMS, EzeScan delivers a fast, cost effective method to transform your manual business processes into intelligent digital workflows. EzeScan benefits include: initiate intelligent automated processes; accelerate document delivery; minimise manual document handling; capture critical information on-the-fly; and ensure standards compliance.

www.ezescan.com.au | info@ezescan.com.au | 1300 393 722

Hyland is a leader in providing software solutions for managing content, processes and cases for organisations across the globe. For 30 years, Hyland has enabled more than 16,000 organisations to digitise their workplaces and fundamentally transform their operations. Hyland has been a leader in the Gartner Magic Quadrant for Content Services for the past 12 years and named one of Fortune's Best Companies to Work For® since 2014, Hyland is widely known as both a great company to work for and a great company to do business with. Our solutions are intuitive to use so organisations can focus on what they do best. Managing information doesn't have to be complicated. At Hyland, our mission is to empower efficiency and agility so our customers can grow and innovate with confidence. We help organisations handle their most critical content and processes with flexible, configurable software solutions.

www.hyland.com/en/| info-onbase@onbase.com| 02 9060 6405

For over 25 years, Informotion's team has specialised in compliance and records management, guiding regulated organisations globally through complexity with clarity, confidence, and proven expertise. We have people in Australia, the UK and Ireland. Today, as data moves to Cloud, AI, and automation, Informotion bridges heritage governance with future-ready innovation. Our practices across Data, Information Governance, Microsoft Cloud & AI are combining decades of compliance mastery with innovative AI, Cloud, and automation tools, to help organisations transform complex information into actionable insights, wherever they operate. Our solutions enable real-time discovery, automated classification, ROT clean-up, compliant retention, and secure management of sensitive information, without compromising compliance. Informotion is a Microsoft Solutions Partner with designations in Data & AI, Digital & App Innovation, Modern Work, Security and Infrastructure. We are the Global Principal Partner for EncompaaS and an OpenText Analytics and Portfolio partner.

www.informotion.com.au| info@informotion.com.au | 1300 474 288

DocuVAN is a Distributor and Reseller of higher end scanning equipment, including Ricoh's state-of the-art scanning solutions in the workgroup, departmental, and production-level scanner categories Ricoh fi Series Best-in-Class Document Scanners deliver speed, image quality, and great paper handling, along with easy integration and compatibility with document imaging applications. We also represent Image Access in Australia, NZ, Pacific Islands and PNG as the distributor of their suite of Bookeye and WideTEK Scanners. If it is deemed part of your core business, Docuvan can supply, install and train you to operate your own scanning solution. We can help you integrate with a document management system and setup workflow processes to automate most paper based legacy systems. Our solutions are scalable and we offer a wide variety of options to suit most budgets.

www.docuvan.com.au| info@docuvan.com.au | 1300 855 839

OPEX® Corporation is a global leader in Next Generation Automation, providing innovative, unique solutions for warehouse, document and mail automation. With a comprehensive suite of customised, scalable technology solutions, OPEX helps clients transform how they conduct business—improving workflow, reducing costs and driving efficiencies in infrastructure. Since 1975, the family-owned and operated company has served as a trusted partner to clients around the world, with nearly 1,600 employees continuously reimagining automation technology that solves the most significant business challenges of today and in the future. OPEX is headquartered in Moorestown, NJ, with facilities in Pennsauken, NJ; Plano, TX; France; Germany; Switzerland; the United Kingdom; and Australia. The year 2025 marks a significant milestone—the company's 50th anniversary under the multi-generational leadership of the Stevens family.

https://opex.com | info@opex.com

Established in 2003, iCognition is a leading Information Management and Governance (IMG) specialist. With over 20 years of customer success stories in delivering IMG services and solutions, we provide managed services for OpenText Content Manager (formerly TRIM) to over 130 government and private sector enterprises across Australia. With information governance at our core, iCognition empowers customers in their digital transformation projects to maximise the value of their information assets. Whether that be on-premises or transitioning to our secure cloud solution, Ingress by iCognition, we enable customers to meet the challenges of managing information across the enterprise. Ingress is a Content Services Platform with OpenText Content Manager at its heart. We can transition your Content Manager system to Ingress or provide a greenfields solution in your cloud or ours. Our Ingress cloud is ISO27001 Information Security Management certified and IRAP assessed to PROTECTED.

www.icognition.com.au | info@icognition.com.au| 1300 4264 00

ELO Digital Office delivers scalable ECM and workflow automation solutions across Australia, New Zealand and the Pacific. Our platform centralises documents, emails and records, helping organisations improve governance, efficiency and collaboration.Key Capabilities:

- Enterprise Content Management & document automation
- Workflow management across all departments
- Records management & compliance (incl. ELO eARC)
- Contract, invoice, HR and learning management modules
- Integration with ERP, CRM, HR and cloud systems

Our services include consulting and solution design, implementation and migration, as well as integration and customisation to meet specific business needs. We also provide comprehensive training and ongoing support to ensure long-term success. ELO's secure, modular and cloud-ready platform scales effortlessly to organisations of all sizes.

www.elodigital.com | info@elodigital.com.au | 1300 066 134

Kapish (a Citadel Edge company), established in 2007, is a dynamic organisation delivering secure technology solutions and strategies in Information Management & Governance, Business Transformation and Enterprise Architecture. Kapish is a Tier 1 OpenText Platinum Business Partner, delivering secure cloud-based information governance and records management solutions built around OpenText's Content Manager (formerly TRIM/HPE RM/MICRO FOCUS CM). Kapish's offerings include IRAP-assessed, ISO 27001-certified cloud managed services, data privacy and protection solutions, IM and technical consulting, migration and implementation services, custom product development and software solutions. Our range of integrated software solutions and managed services gives you a complete view of your IT landscape, helping you discover, manage and protect your information assets, meet regulatory compliance, boost user productivity and transform business processes with modern solutions.

kapish.com.au | info@kapish.com.au | 03 9017 4943

Newgen offers a unified digital transformation platform that includes native process automation, content services, and communication management capabilities. Globally, many successful enterprises across various industries rely on the NewgenONE digital transformation platform—a comprehensive and unified cloud-based platform with low code capability for rapid development of content-driven, customer-engaging business applications. The platform can transform and simplify complex business processes. Equipped with cutting-edge technologies, including mobility, social listening/sensing, analytics, cloud, artificial intelligence (AI), machine learning (ML), and robotic process automation (RPA), the NewgenONE platform helps enterprises stay ahead of the curve. From grass-root citizen experience management, dynamic case management to electronic documents and records management, lending to underwriting, the platform solves multiple use cases across various industries, including government, banking, insurance, and others.

newgensoft.com | info@newgensoft.com | 02 80466880

NZ Public Sector Embraces Document Automation

Wellington City Council and Fire and Emergency New Zealand (FENZ) have deployed intelligent document processing systems to meet new regulatory requirements, eliminating manual data entry for rates rebates and invoice processing.

The implementations address urgent compliance pressures facing New Zealand's public sector.

Wellington City Council faced regulatory mandates from the Department of Internal Affairs requiring automation of rates rebate processing by early 2025.

FENZ needed to comply with new payment time reporting requirements while managing invoices across more than 650 stations.

Both organisations deployed ABBYY Document AI technology through integration partner Desktop Imaging.

The vendor claims the systems achieved "100% compliance" and "zero manual data entry" for Wellington's rates rebate processing.

Absolutely Positively Wellington City Council
Me Heke Ki Pōneke

Wellington City Council

Following new regulatory mandates from the Department of Internal Affairs (DIA) in early 2025, Wellington City Council faced an urgent need to modernize its manual rates rebate process.

By implementing ABBYY Document AI through Desktop Imaging's DI Automation: Rates Rebate Service, the Council fully automated form validation, eligibility checks, reporting, and document retention.

More than 800 applications have been processed with zero manual data entry, achieving 100% compliance.

Following comprehensive planning and testing, the actual deployment of the cloud-based solution took just days, with no additional IT infrastructure required.

"By using purpose-built AI for intelligent automation, Wellington City Council have eliminated manual bottlenecks and ensured every application is processed accurately and on time," Beighton said.

"Our expertise in New Zealand local government processes was invaluable in customising the solution to meet Wellington City Council's specific requirements and Department of Internal Affairs compliance standards."

WHAKARATONGA IWI
FIRE EMERGENCY
NEW ZEALAND

Fire and Emergency New Zealand

Fire and Emergency New Zealand (FENZ) processes tens of thousands of invoices annually across more than 650 stations.

Manual handling and legacy systems caused delays, errors, and compliance risks.

ABBYY Vantage is deployed as part of Desktop Imaging's DI Invoicing cloud service to automatically capture, classify, and extract the data from invoices to feed the DI Invoicing validation engine for business rule validation, fraud checks, and IRD-compliant workflows.

The result is dramatically reduced manual processing, real-time compliance, faster payments, and scalable automation without additional infrastructure or staff.

"Improving accuracy and speed were top priorities for FENZ as we needed to adhere to new payment time reporting requirements. But we also wanted an 'off-the-shelf' solution that could be deployed quickly without adding complexity," said Kris Elliott, Solutions Specialist at Desktop Imaging.

"By deploying ABBYY technology as an integral part of our DI Invoicing service for customers like FENZ, we improved accuracy, data quality, compliance, and speed.

"The outcomes have been immediate — faster payments, greater transparency, and more time for their teams to focus on supporting the communities they serve."

The implementations reflect accelerating adoption of intelligent document processing across New Zealand's public sector.

Regulatory mandates from central government agencies are driving automation investments at both national and local levels.

https://www.abbyy.com

https://www.desktopimaging.co.nz

Vantage 3.0 adds LLM integration

Document automation vendor ABBYY has launched Vantage 3.0, with the Document AI platform now offering direct integration with LLMs

It also introduces a powerful analytics dashboard, built-in sophisticated compliance and redaction capabilities and an expanded library of pre-configured AI models that enable organizations to gain real-time insights, protect sensitive information, and accelerate process transformation.

The company says the combination of its purpose-built AI with the flexibility of GenAI allows users to leverage the vast capabilities of LLMs safely within business process workflows.

"This hybrid approach ensures users expand document automation capabilities while maintaining the control, consistency and explainability required for enterprise-grade solutions."

ABBYY Vantage seamlessly connects to Azure OpenAI for prompt-based extraction without complex custom coding. It includes pre-engineered prompts optimized for data extraction to reduce setup time, and validation to ensure consistent results and continuously improve performance.

"We know that context matters when it comes to data extraction. With ABBYY Vantage, you can control exactly how data is sent to the LLM and choose between sending the document image or the structured, precise text output from ABBYY's highly acclaimed OCR," commented Max Vermeir, Vice President of AI Strategy at ABBYY.

"Furthermore, you can track where data was extracted from within the document for auditability and transparency that's essential for compliance."

Vantage 3.0 embeds robust compliance controls, including enterprise-grade redaction tools that automatically remove sensitive data before storage or export. This ensures healthcare, financial services and insurance organizations safeguard confidential information, meet regulatory demands, and reduce the risk of data breaches.

Enhanced role-based access controls and rigorous audit trails further fortify compliance for mission-critical document processing providing peace of mind for privacy-conscious enterprises.

Furthermore, businesses running mission-critical document processing benefit from strengthened BC/DR capabilities that ensure uninterrupted operations, even in the face of regional or infrastructure failures.

ABBYY Vantage adheres to SOC 2, ISO certifications, GDPR, CCPA, FIPS, and STIG compliance.

A new comprehensive analytics dashboard provides organizations with granular process insights and realtime performance metrics such as touchless processing rates, document type detection accuracy, human-in-the-loop (HITL) corrections.

https://www.abbyy.com

Automated Workflow for AI agents

Elastic is targeting organisations seeking to automate complex workflows using AI agents that can access and analyse enterprise data with its new Agent Builder.

The platform addresses the challenge of deploying AI systems that can reliably interact with internal documents, databases and systems while maintaining governance and security controls.

Agent Builder connects to Elasticsearch to retrieve enterprise data and execute automated tasks.

The company claims the system "dramatically simplifies" agent development through integrated data ingestion, retrieval capabilities and built-in tools.

Developers can deploy agents using multiple AI models through cloud providers.

Elastic simultaneously introduced Workflows, currently in technical preview, which adds rule-based automation to supplement AI-driven decision-making.

This dual approach responds to enterprise requirements for both intelligent reasoning and predictable, compliant automation – a critical consideration for organisations managing regulated data.

"Agent Builder has native MCP and A2A protocol support, enabling seamless deployments within Microsoft Foundry and Microsoft Agent Framework," said Amanda Silver, corporate vice president at Microsoft CoreAI.

"This gives our users a way to build context-rich, agentic AI leveraging Elasticsearch as a Knowledge Source."

Elastic has partnered with Arcade.dev for agent tool integration and LlamaIndex for document processing.

"Unlocking enterprise context from unstructured data sources is key to building effective agents," said Jerry Liu, CEO at LlamaIndex.

The company claims its document processing capabilities will help agents "retrieve, process, and prepare data so agents can reason more accurately."

Agent Builder is available in Elastic Cloud Serverless and included with the Enterprise Tier for existing customers using Elastic Cloud Hosted or self-managed deployments.

Ken Exner, chief product officer at Elastic, positioned the release as addressing a gap between AI experimentation and production deployment: "By enhancing Agent Builder with Workflows, teams get a single system that delivers both intelligent reasoning and dependable automation, which is exactly what enterprises need to move from pilots to real-world impact."

https://www.elastic.co

Wondershare upgrades KM tool with AI

Wondershare has released EdrawMind V13, adding node-based note capabilities to its mind mapping software aimed at knowledge management workflows.

The update introduces a dual-pane interface where users can create mind maps alongside detailed notes. The vendor claims this is among the first professional tools to support node-based notes.

Notes can include text, images, videos, tables, links and code blocks. The system displays the mind map's structural overview on the left while detailed context appears on the right.

The release adds AI capabilities including a webpage summariser that converts web content into mind maps. The AI summariser generates notes for key nodes, according to Wondershare.

An AI web search function creates mind maps and presentations with citation sources. The software also supports flowcharts, relationship line drawing and note sketching.

"By combining AI intelligence with intuitive design, we're empowering users to transform information into insight faster," said Aiden, Head of Wondershare EdrawMind.

The interface redesign includes a left-panel entry for AI-generated mind maps and eight thematic styles.

https://www.edrawmind.com

Process designer converts text into visual workflows

Canadian knowledge management vendor Procedureflow has released two AI-powered features targeting organisations managing complex compliance processes and procedural documentation.

The company's AI Search tool uses natural language processing to interpret search queries, while AI Process Designer converts text-based procedures into editable visual workflow diagrams, according to the vendor.

Procedureflow CEO Daniella DeGrace said the tools aim to reduce the time required for agents to find procedural information and for process owners to create documentation.

The AI Search feature analyses user intent rather than keyword matching, the vendor claims. Users can enter conversational queries to retrieve relevant process workflows from the knowledge base.

AI Process Designer uses generative AI to transform existing text documents into structured flowcharts. Users can import procedures and edit the resulting diagrams, though the company did not specify which document formats are supported or integration capabilities with existing systems.

The features address challenges identified in knowledge management research. Gartner estimates poor data quality costs organisations an average of $US12.9 million annually, while the research firm notes AI and machine learning technologies make foundational knowledge management practices more essential rather than eliminating them.

https://www.procedureflow.com

Kodak Alaris Expands IDP Platform With GenAI

Kodak Alaris has extended its Info Input Solution IDP platform beyond traditional document AI to include generative AI services from Google, OpenAI, AWS, and Box.

Version 7.5 introduces native integrations with Google Gemini, AWS Bedrock Data Automation, ChatGPT, and BoxAI. These join existing connections to Google Doc AI, Microsoft Document Intelligence, Hyperscience, and Amazon Textract.

The update addresses organisations managing complex compliance processes and digital transformation initiatives. New features include a drag-and-drop workflow editor and integration of the IRIS OCR engine at no additional cost.

The expansion into generative AI represents a strategic shift for document processing platforms. While traditional document AI excels at data extraction and classification, generative AI can summarise multi-page documents and provide contextual analysis.

Kodak Alaris describes its approach as "Open Intelligence design", a vendor claim that emphasises integration flexibility over proprietary AI models. The platform allows organisations to connect multiple AI services within a single workflow.

"Our mission is to make document intelligence open, application-oriented, and sustainable for the long-term," said Megan Bevilacqua, Senior Product Manager for Kodak Alaris.

"Info Input Solution 7.5 does exactly that - we reduce complexity, create more room for innovation, and enable customers to implement automation faster and continuously develop their processes in line with advances in AI."

The IDP platform can handle any document from any source, including unstructured or handwritten files, and automatically verify critical details such as legal references, threshold amounts, and content plausibility.

Automatic summarization provides a quick overview of multi-page documents or email threads, helping employees classify processes faster and save time on recurring routine tasks.

https://www.kodakalaris.com

IBM Bets Big on Real-time Data Streaming

IBM has acquired data streaming platform provider Confluent for $US11 billion. Confluent operates a commercial distribution of Apache Kafka, an open-source platform that enables realtime data streaming between applications and systems.

IBM claims the transaction addresses growing enterprise requirements for real-time data integration as organisations deploy generative AI and manage increasingly complex information architectures.

Data streaming refers to the continuous flow of information between systems as events occur, rather than processing data in batches.

Traditional enterprise architectures move data in scheduled intervals - nightly batch jobs transfer records between databases, for example.

Data streaming creates persistent connections that transmit information the moment it is generated.

Apache Kafka functions as a distributed message broker that captures, stores and routes data streams.

Applications publish events to Kafka topics - categorised feeds of information.

Other applications subscribe to these topics and receive events in real time.

Kafka retains event logs, enabling systems to replay historical data or recover from failures.

It addresses a fundamental challenge: maintaining data consistency across siloed systems.

When customer records update in a CRM system, inventory changes in an ERP platform, or transactions process through payment systems, streaming platforms propagate those changes immediately to downstream applications, data warehouses and analytics tools.

The technology has direct implications for records management and compliance.

Data streaming platforms create immutable event logs - permanent records of every state change across systems.

This provides audit trails showing exactly when information was created, modified or accessed across the enterprise.

Organisations managing complex compliance requirements use streaming platforms to enforce data governance policies in realtime.

When personal information moves between systems, streaming infrastructure can apply masking, encryption or access controls at the point of transfer rather than relying on downstream systems to enforce policies.

"IBM and Confluent together will enable enterprises to deploy generative and agentic AI better and faster by providing trusted communication and data flow between environments, applications and APIs," said Arvind Krishna, IBM chairman, president and chief executive officer.

Technical Architecture

Confluent's platform extends Apache Kafka with enterprise features including pre-built connectors for popular enterprise systems, stream processing capabilities, and governance tools.

The platform integrates with major cloud platforms including AWS, Google Cloud Platform and Microsoft Azure, as well as data warehouses like Snowflake.

Deployment options include fully managed cloud services, self-managed on-premises installations, and hybrid bring-your-own-cloud configurations.

The platform serves more than 6,500 clients. The company states that more than 40% of Fortune 500 organisations use its services.

Jay Kreps, Confluent's CEO and co-founder, co-created Apache Kafka while working at LinkedIn before founding Confluent in 2014.

Six Principles for Building AI Agents That Work

Skan AI has released a new framework addressing the critical gap between AI capabilities and enterprise process automation needs.

The "Agentic Process Automation Manifesto" establishes six principles for developing effective AI agents based on real-world process telemetry rather than theoretical workflows.

The manifesto draws from more than 50 deployments at major enterprises in banking, healthcare, insurance, and services sectors where traditional automation approaches have fallen short.

"Every large enterprise wants AI that can reason and act. The blocker is that agents lack an accurate picture of how work is actually performed," said Manish Garg, Co-founder and Chief Product Officer of Skan AI.

Central to the framework is Skan's Observation-To-Agent (O2A) platform, which creates a "living blueprint" of work processes by capturing human-system interactions across enterprise applications.

The company identifies four current automation approaches it claims consistently fail: task-first approaches that lack flexibility, database-first methods that provide insight without execution, platform-constrained systems, and models trained primarily on consumer rather than enterprise patterns.

In contrast, Skan's manifesto outlines six core principles:

1. Telemetry over assumptions - Training on observed interactions across approved applications

2. Execution over analytics - Enabling end-to-end execution rather than just insights

3. Transparent governance - Building in policy-as-code, case memory, and audit trails

4. Open architecture - Integrating with existing systems without replacement

5. Outcome-driven metrics - Measuring business outcomes rather than automation volume

6. Human-AI collaboration - Establishing clear roles between AI and human workers

The approach promises more consistent adherence to policy requirements while maintaining complete documentation of processes.

"When operational telemetry informs agents and controls are built in from the start, you see faster resolution times and more consistent compliance," said Cijo Joseph, Chief Technology & Digital Officer at Mitie.

The complete manifesto is available at https://www.skan.ai.

Law firm automates claims management

Wotton Kearney will replace legacy systems with an AI-powered platform to manage legal matters and client reporting. The Asia-Pacific insurance and risk legal firm has partnered with process automation vendor Appian to deploy a system called OSCAR (Organised System for Claims and Reporting).

The system is expected to go live by April 2026 and will handle automated data extraction, email classification and document generation.

Wotton Kearney said the platform will save 'significant lawyer hours annually'.

The 500-lawyer firm currently relies on manual processes for client reporting and matter tracking that 'consume substantial resources', according to the announcement.

Built on the Appian Platform, OSCAR will use intelligent document processing and integrate with Wotton Kearney's financial ERP system.

'By integrating Appian, we're enhancing accessibility, responsiveness and value for our clients,' said Charles Simon, managing partner for casualty and operations at Wotton Kearney.

Australian law firms are racing to adopt AI tools amid growing competition and client expectations for digital capabilities.

Major firms including Clayton Utz, MinterEllison and Herbert Smith Freehills have trialled generative AI for legal research and document drafting. A December 2024 survey found 65% of law firms now have an AI strategy or responsible use policy.

The deployment comes as Australian organisations face heightened data governance requirements under Privacy Act reforms that commenced in 2024 and continue through 2025.

Wotton Kearney operates across Australia, New Zealand, Singapore and Thailand with 95 partners and over 500 lawyers.

ECM Vendor Rebrands as Doxis

German enterprise content management (ECM) vendor SER Group has rebranded as Doxis. The Bonn-based company, which has announced plans to launch into the Australian market, now positions itself as "The Document Intelligence Company" following three acquisitions in 2025.

The rebrand aligns the company name with its flagship Intelligent Content Automation platform, which the vendor claims is used by five million people daily. Product and service names remain unchanged.

Yeelen Knegtering, co-founder and CEO of intelligent document processing vendor Klippa, joins as Chief AI Officer. Doxis acquired Klippa in March 2025. Knegtering will oversee AI strategy, implementation, ethics and governance across the business.

CEO Dr John Bates said the rebrand followed profitable growth in 2025 and three strategic acquisitions.

The vendor claims its unified platform differentiates it from competitors offering "fragmented point solutions."

The Doxis platform integrates with SAP, Salesforce and Microsoft enterprise systems. The rebrand reflects broader consolidation trends in the enterprise content management and intelligent document processing markets. Vendors are increasingly positioning AI-powered document automation as a unified capability rather than standalone tools.

https://www.doxis.com

OpenText Brings Capture Software to Private Cloud

OpenText has added a private cloud deployment option for its Capture document processing platform. The single-tenant environment aims to address organisations requiring cloud scalability while maintaining data sovereignty and compliance control.

The offering sits within OpenText's Private Cloud infrastructure, where the vendor manages hosting, maintenance and operations. Customers retain control over upgrade timing and access to design and administration tools.

OpenText claims the deployment model can reduce IT costs by 30-40% through elimination of physical server maintenance and unused resource charges. The platform uses AI to automate document classification, data extraction and validation across multiple input channels including bulk scanning, email, fax and mobile documents.

OpenText positions this as addressing high-volume use cases such as invoice processing, HR onboarding and customer service workflows.

With the latest updates like OpenText Capture 25.4, business administration teams gain web access to intuitive dashboards for real-time monitoring and instant data corrections, enabling them to stay in control of capture operations from any device.

Private cloud deployments represent a middle path between public cloud and on-premises infrastructure. Organisations in regulated industries increasingly seek this model to balance modernisation requirements with data residency and compliance obligations.

The division of responsibilities places infrastructure management, security patching and performance monitoring with OpenText, while customers configure workflows and business processes.

https://www.opentext.com

MediRecords targets inbox overload

Medical practices could save up to two hours daily managing incoming documents under an AI automation tool entering early release in 2026.

MediRecords' Evolve Direct automates document classification, patient matching and provider assignment for clinical inbox management. The vendor claims the tool reduces manual processing from 18 clicks to three, an 83% reduction.

Beta testing with clinics operating up to 10 doctors showed potential time savings of 120 minutes per day, according to the vendor. This represents approximately 480 hours annually.

The system analyses incoming documents, assigns them to the correct provider, matches them to patient records and performs virus scanning. Unmatched documents are directed to a holding area for manual review.

Healthcare administrative burden has intensified as practices manage increasing document volumes from pathology results, specialist reports and patient correspondence.

Staff traditionally download documents from external systems, locate the relevant patient file and manually upload attachments, a workflow vulnerable to misfiling errors. MediRecords founder and chief executive Matthew Galetto said the tool addresses cost and time pressures.

"Evolve Direct provides a significant boost for clinics under increasing cost and time pressures. By largely automating inbox management, we're helping practices reclaim hours each week, reduce risk and focus on delivering quality care," he said.

The release forms part of a broader multi-agent AI strategy. MediRecords has deployed Evolve Patient Summary, which generates clinical record snapshots, with additional AI agents planned for 2026.

https://www.medirecords.com

AI adopted for Payroll Compliance

More than three-quarters of Australian organisations now use artificial intelligence to manage payroll compliance, according to new research examining responses to wage theft reforms.

The 2026 State of Payroll Compliance Report reveals 77% of employers deploy AI to detect compliance issues, track legislative changes and review contracts.

The most common application is monitoring payroll to detect compliance issues, used by 42% of organisations surveyed.

Another 57% plan to implement automated payroll review technology, reflecting growing recognition of technology's role in managing complex compliance requirements.

However, the report cautions that technology alone cannot ensure compliance without human oversight for interpretation and validation.

"AI has become an enabler in payroll compliance," the report states. "It can analyse large volumes of data faster and make sense of complexity at scale."

"But technology alone is not assurance," it continues. "The human layer which involves interpreting, validating, and applying judgment is what turns insight into confidence."

The findings follow January 2025 wage theft reforms that reshaped employer approaches to payroll accuracy and governance.

Almost nine in 10 organisations (89%) introduced new payroll compliance measures after the reforms, including staff training, audits and system improvements.

Despite heightened focus, challenges persist. Only 64% of organisations report full confidence in their payroll compliance, while 36% remain unsure they pay employees correctly.

The top compliance concern is interpreting employment instruments such as modern awards and enterprise agreements, cited by 43% of respondents.

System limitations handling unique or complex scenarios ranked second at 40%, followed by poor integration between systems.

Marcus Zeltzer, founder and managing director of Yellow Canary, which commissioned the research, said most payroll issues are unintentional.

"The most significant payroll issues are rarely intentional," Zeltzer said. "Instead, they are driven by the complexity of modern award interpretation and the limitations of payroll systems."

The report found organisations conducting regular audits show significantly higher confidence levels than those auditing only when issues arise.

Quarterly reviews are most common (35%), while only 5% conduct reactive audits. Organisations using reactive approaches report notably lower confidence.

Manual spreadsheet audits declined from 31% to 23% as employers shift toward technology-driven assurance methods.

Average spending on payroll compliance audits reached $272,000 in 2025, though organisations allocated $507,227 on average, suggesting budget under-utilisation.

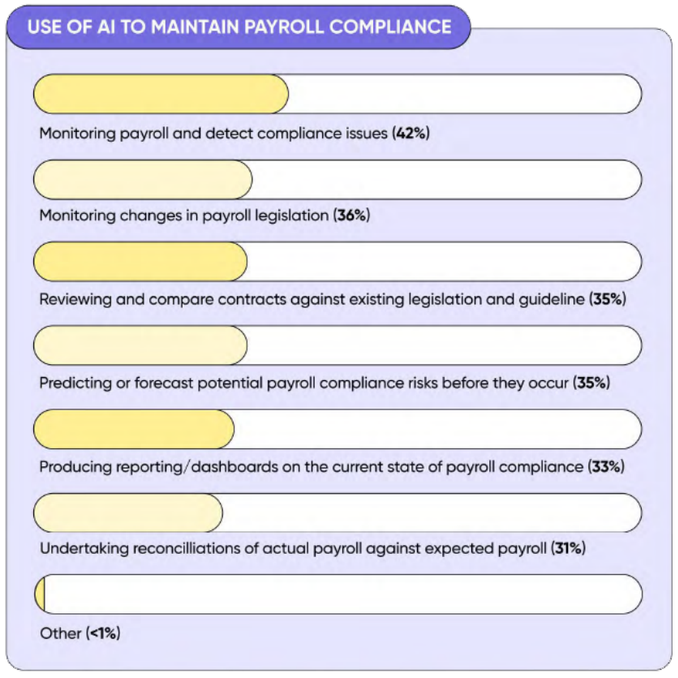

Larger employers with 5,000-plus employees spent an average of $434,439, compared with $186,415 for organisations with 50-199 employees.

Looking ahead, one-third of organisations plan to strengthen all compliance areas collectively in 2026.

Payroll tax (22%) and awards/enterprise agreements (19%) emerged as the most pressing individual priorities for improvement.

The Fair Work Ombudsman recovered $358 million for over 249,000 underpaid workers in 2024-25, highlighting the scale of payroll errors.

"Payroll underpayments can trigger a snowball effect," the report warns. "One error can uncover misclassifications, award mistakes, super, leave and payroll tax issues."

The research was conducted by Lonergan Research between 20-24 October 2025, surveying 540 Australian business and payroll decision-makers from organisations with 50-plus employees.

https://www.yellowcanary.com.au/

FrankieOne powers PEXA compliance

Property settlement platform operator PEXA Group is entering the compliance software market ahead of mandatory anti-money laundering regulations for real estate professionals taking effect in July.

The company's PEXA Clear solution will launch the same month AUSTRAC's Tranche 2 AML/CTF regulations commence, extending anti-money laundering obligations to real estate agencies, conveyancers and legal practitioners for the first time.

AUSTRAC's Tranche 2 regulations represent a significant expansion of anti-money laundering obligations into the property sector. The regulations were developed following concerns about property transactions being used to launder criminal proceeds, a long-standing issue highlighted by organisations including Transparency International and the Financial Action Task Force.

The timing presents both a market opportunity and operational challenge. With registrations opening 31 March 2026, practitioners will have less than four months to implement and test compliance systems before the regulatory deadline.

PEXA Clear is being developed in partnership with FrankieOne, a global identity verification platform that counts Commonwealth Bank, Westpac and Shopify among its clients. The Real Estate Institute of Queensland will provide training support for real estate agents.

The solution handles identity verification using selfies and government-issued IDs, conducts Know Your Business and Know Your Customer checks, and screens individuals against Politically Exposed Persons lists, sanctions databases and adverse media sources.

PEXA claims the system provides "actionable guidance" on when to escalate risks, including prompts for Suspicious Matter Reports to AUSTRAC. It maintains audit trails to support compliance reporting requirements.

The company describes PEXA Clear as "purpose-built" for the Australian property industry, though it has not disclosed pricing, expected customer numbers, or specific technical integration details with existing property management and legal practice systems.

Amelia Howell, Head of Customer Growth and Solutions at PEXA Clear, said the company's decade-long involvement in Australian property settlements provides insight into transaction risks.

"PEXA Group has a comprehensive view of not just the complexities, but also the touchpoints for risk in a property transaction," she said.

FrankieOne CEO Simon Costello said the partnership "combines FrankieOne's best-of-breed identity, KYB and risk orchestration capabilities with PEXA Clear's deep understanding of property workflows."

The solution will operate independently from PEXA's Electronic Lodgement Network for property settlements, which has facilitated more than 20 million transactions since 2013.

Product demonstrations are currently available, with a series of roadshows scheduled to commence in February.

https://www.pexa.com.au

IBM Unveils Agentic AI breakthrough

An eight-week proof of concept has tested whether autonomous AI agents can operate within enterprise compliance frameworks while meeting governance requirements. IBM and global technology firm e& completed the trial using IBM watsonx Orchestrate integrated with the OpenPages governance platform.

The companies claim the technology represents a shift from question-and-answer chatbots towards AI systems that can take autonomous actions. The system is designed to help employees and auditors access and interpret legal, regulatory and compliance information with what IBM describes as "clear, traceable responses aligned with enterprise governance requirements".

IBM claims the deployment represents "one of the early enterprise-grade agentic AI implementations in the region". The companies announced the collaboration at the World Economic Forum in Davos.

IBM's watsonx Orchestrate provides access to more than 500 tools and customisable agents from IBM and partners. The system's architecture allows large language models to run across hybrid environments, including customer-managed infrastructure, while maintaining governance controls. This addresses concerns about data sovereignty and regulatory requirements that affect many organisations considering AI deployment.

The proof of concept tested whether AI agents could reason, orchestrate tasks and integrate with enterprise systems while remaining aligned with e&'s existing governance, risk and compliance framework. Success depends on the technology's ability to provide explainable decisions that auditors and regulators can verify.

Hatem Dowidar, Group CEO of e&, said the company aims to move "beyond isolated AI use cases toward enterprise-scale agentic AI that is trusted, governed, and deeply integrated into how the organisation operates".

e& is a global technology and investment conglomerate based in Abu Dhabi, United Arab Emirates

Ana Paula Assis, IBM's SVP and Chair for Europe, the Middle East, Africa and Asia Pacific, said the proof of concept "intends to demonstrate how agentic AI can be designed and validated for enterprise-scale use, deeply integrated into core systems, governed by design, and trusted to support human-led decisions and outcomes".

UiPath Acquires WorkFusion Agents

Automation vendor UiPath has acquired WorkFusion in an undisclosed deal to strengthen its portfolio of AI-powered compliance solutions for financial services.

WorkFusion specialises in AI agents for financial crime compliance, including anti-money laundering (AML), know your customer (KYC), sanctions screening and transaction monitoring investigations.

The company, which describes itself as a "pioneer" in the field, provides pre-built AI agents that automate labour-intensive compliance workflows.

The acquisition addresses mounting pressure on financial institutions to manage increasingly complex compliance requirements while reducing operational costs. Banks globally face growing regulatory scrutiny over financial crime prevention, with compliance teams struggling to process high volumes of alerts and investigations manually.

"Joining UiPath is a moment of validation for the years our team has poured into creating something bold, different and deeply needed in financial crime compliance, AI agents that automate work and mitigate risk," said Adam Famularo, CEO of WorkFusion.

"UiPath gives us the scale to grow faster than we ever could alone."

UiPath CEO Daniel Dines said financial institutions need intelligent solutions to combat sophisticated financial crimes and navigate evolving compliance requirements. "Incorporating WorkFusion's purpose-built AI agents for financial crime compliance into our platform expands our portfolio of agentic AI solutions for these industries," he said.

The combined offering will enable banks to automate complex compliance workflows, analyse patterns and prioritise cases requiring human expertise while maintaining security, governance and regulatory controls, according to UiPath.

The deal strengthens UiPath's position in the growing market for compliance automation as organisations seek to balance regulatory obligations with operational efficiency. Automation of compliance processes represents a significant opportunity for enterprise technology vendors targeting financial services, where manual review processes remain resource-intensive and error-prone.

https://www.uipath.com

AI agents beat RPA in unstructured data

AI-powered document processing agents achieved 40% higher accuracy than robotic process automation systems when handling unstructured documents, according to a study released by Artificio Products Inc.

The vendor-commissioned research analysed over 500,000 document processing transactions across healthcare, finance, real estate and logistics sectors. Traditional RPA systems performed adequately with standardised documents but struggled with variable formats, the study found.

AI agents achieved 94% accuracy on medical forms with variable layouts compared to 61% for RPA systems, according to the study. Financial institutions processing loan documentation reported 89% straight-through processing with AI agents versus 53% with RPA.

"AI agents don't just match patterns, they comprehend document structure and meaning, handling real-world variability that breaks traditional automation," said Lal Singh, founder and CEO of Artificio Products Inc.

The study tracked field-level extraction accuracy, exception rates, processing speed and adaptability. AI agents reduced exception handling time by 67% compared to RPA implementations, the research claimed.

Artificio's document intelligence agent identified document types with 97% accuracy without templates, while RPA required pre-configured specifications. The vendor's multi-agent architecture coordinates specialised agents from intake through ERP integration.

Organisations implementing AI agent technology reported 80% reduction in configuration and maintenance costs compared to RPA systems, according to the study.

AI agent implementations delivered ROI 3.2 times higher than RPA deployments over three years, primarily due to reduced manual intervention.

The platform processes documents through pipelines combining computer vision, natural language processing and custom-trained models. AI agent deployments achieved production readiness within four to six weeks, compared to six to nine months for RPA implementations.

The comparative analysis examined processing performance across 500,000+ documents from 47 enterprise implementations over 18 months, according to Artificio.

https://artificio.ai

PaperCut one-click data extraction

Intelligent document processing capability has been embedded directly within PaperCut MF and PaperCut Hive. AIDA, an AI-powered document processing platform, is partnering with Selectec to integrate advanced data extraction and automation features into PaperCut's scanning and print workflows.

The integration enables organisations to move beyond standard optical character recognition, immediately extracting data from scanned documents without templates or manual configuration.

This integration also enables data input at the MFP. Teams can append metadata directly from the touchscreen, streamlining workflows like expense tracking or invoice routing without leaving the device.

The partnership addresses a critical gap in enterprise document workflows. Organisations managing complex compliance processes face challenges balancing automation with governance requirements.

Compliance managers, records managers and information governance teams must ensure documents are properly classified, validated and routed according to regulatory standards whilst reducing manual processing time.

The AIDA platform leverages hybrid artificial intelligence technology, allowing the system to learn from a single document example. The technology processes both structured and unstructured documents including invoices, contracts, forms and receipts without requiring predefined templates or extensive training data.

Once data is extracted, organisations can trigger automated workflows including archiving, task distribution, payment processing or business intelligence reporting.

This capability integrates with PaperCut's existing governance features. PaperCut MF provides comprehensive audit trails, job logging and document archiving functionality critical for regulated industries including healthcare, finance and legal sectors.

PaperCut's recent 24.1 release introduced enhanced scan reporting, enabling visibility into who scanned what, where and when, with custom filtering capabilities for compliance monitoring.

https://www.papercut.com

AI Security Framework from OpenText

OpenText has unveiled enhanced security measures to protect AI implementations across enterprises by unifying defences across identity management, data protection, application security and operational workflows.

The new capabilities announced in Cloud Editions 25.4 include Core Identity Foundation for unified access controls, Application Security Aviator with AI-powered auto-remediation, and expanded threat detection services with built-in compliance controls.

"To keep pace with AI's velocity, we are giving security and IT teams the tools to work faster, smarter, and with greater confidence," said Muhi Majzoub, OpenText EVP, Security Products.

The Core Identity Foundation aims to address hybrid environment challenges by unifying identity and access management across on-premises, cloud, and legacy systems without requiring infrastructure overhauls. The company describes this SaaS-based solution as implementing Zero Trust controls through an Identity-as-a-Service approach.

For application security teams, the Application Security Aviator 25.4 promises to reduce vulnerability remediation times through automated, validated code fixes integrated directly into DevSecOps workflows via the Fortify Command Line Interface.

OpenText has also enhanced its threat detection capabilities with behavioural analytics and introduced advanced encryption services to protect sensitive data across its lifecycle, including when used for AI training and operations.

The company is offering complementary security assessments to identify redundant, obsolete, and trivial data that could create security vulnerabilities.

https://www.opentext.com

Guidewire Bolsters KM with Pronavigator

Guidewire has entered into a definitive agreement to acquire ProNavigator, an AI-powered knowledge management platform specialised for property and casualty insurers.

The acquisition will integrate ProNavigator's context-aware intelligence across Guidewire's platform, enhancing decision-making capabilities for underwriters, claims adjusters and customer service representatives.

ProNavigator's technology delivers precise information within insurance workflows, helping professionals access institutional knowledge across claims, underwriting and distribution channels.

"ProNavigator, inside of Guidewire, makes every insurance professional an expert," said Mike Rosenbaum, Chief Executive Officer, Guidewire.

Currently used by 34 insurance organisations, including 12 shared customers with Guidewire, the platform helps reduce claim cycle times and accelerate employee onboarding according to the companies.

The integration aims to provide advanced search and contextual knowledge retrieval capabilities that Guidewire claims will help insurers deliver faster responses and more personalised guidance to customers.

ProNavigator CEO Joseph D'Souza and his team will join Guidewire once the acquisition closes, which is expected in the second quarter of Guidewire's fiscal year 2026.

Laura Drabik, Chief Evangelist at Guidewire, described context-aware knowledge as "the next frontier of intelligence for the P&C industry."

ProNavigator was previously a participant in Guidewire's Insurtech Vanguards incubator program before the acquisition.

The financial terms of the acquisition were not disclosed in the announcement.

www.guidewire.com