idm.
information & data manager
JUNE-JULY 2026

Rethinking the Essential Eight

The Hidden Knowledge Crisis Hurting Your Organisation

When Data Moves but Meaning Disappears - how AS5393 Can Help

Do you Believe in the Data Fairies?

NZ Council seeks M365 Records Overhaul

Christchurch City Council has issued a Request for Information (RFI) seeking vendors to deliver a Microsoft 365-based Electronic Document and Records Management System (EDRMS) to replace its existing on-premise TRIM platform. The council is the second largest local government authority in New Zealand with over 3000 staff.

The tender document specifies that the replacement solution must be "based on Microsoft 365" and leverage its ecosystem, explicitly citing SharePoint Online, Microsoft Purview, and Power Platform as core components.

The Council also expects Microsoft Purview to form "a major component of the governance, compliance, and insider risk controls," including monitoring for anomalous user behaviour, privileged access misuse, and data exfiltration prevention. It describes the programme as a multi-year transformation. "The goal is to replace the legacy on-premise system with a modern, compliant, scalable solution based on Microsoft 365 that aligns with statutory obligations and strategic goals," the RFI states. The Council expects the solution to support scalable architecture for 10-plus years of projected growth.

The Records Management Policy cited in the tender documents identifies systemic record-keeping failures driving the project. "Council records are not consistently saved in the Council's records management system (TRIM)," the policy states. Poor document naming, information stored on individual PCs, flash drives and email inboxes, and hard-copy records not returned to storage are all flagged as systemic problems.

The policy warns these shortcomings lead to "less than ideal decision-making based on incomplete information and difficulties discovering the correct information when responding" to information requests.

Automated Metadata and Classification

A central objective of the tender is to automate metadata capture and record classification at scale. The functional requirements state that "record identification, classification, and retention should be automated as much as possible, minimizing human intervention to only exceptional cases."

The solution must automatically scan both on-premise and cloud repositories, identify and classify records based on file metadata and content, and assign sensitivity labels - such as In-Confidence, Sensitive, or Unclassified - based on content analysis. It must also detect duplicate records across repositories and identify critical records.

Automated capture of the minimum metadata fields mandated under the Archives NZ Information and Records Management Standard is a formal requirement. These fields include: unique identifier, name and title, date created, business activity, creator, software version, and a full log of any later actions carried out on the record, including the persons or systems involved and dates.

The tender specifies that manually editing system-generated metadata fields must be restricted to authorised users only, to preserve chain-of-custody integrity.

The Council requires full automation of its retention and disposal lifecycle. The system must automatically identify records that have reached their retention period and flag them for archiving or destruction in accordance with the retention schedule. Destruction of records requires approval from authorised users in the Information Management team, with the tender specifying the process must be "defensible, auditable, and compliant with Archives New Zealand standards."

Enterprise System Integrations

The tender identifies an extensive list of Council enterprise systems the new EDRMS must integrate with. The Council's current collaboration environment - Microsoft Teams, SharePoint Online, and OneDrive - forms the core daily work platform. Outlook and Exchange Online are flagged as critical for email records management, eDiscovery, and LGOIMA compliance, requiring retention and legal hold policy support.

The tender also specifies integration requirements with a range of systems including SAP – ERP, ArcGIS, and Pathway.

The new EDRMS must support governance, discoverability, integration, and lifecycle management of these assets, including upload, preview, retrieval, streaming, and lifecycle governance of high-resolution images and large video files without performance degradation.

An AI-ready architecture is specified as a future-readiness requirement, supporting automated classification, semantic search, and predictive analytics. The tender states new capabilities such as AI utilisation are in the Council's future roadmap, with the current programme focused on delivering incremental value through a phased rollout.

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

Your Organisation Has a Knowledge Management Problem. It Just Doesn't Know It Yet

By Stephanie Barnes

Ask most people what Knowledge Management means and they'll mention a chatbot. Or a FAQ database. Or the system the contact centre uses to answer customer calls faster.

They're not wrong. That *is* KM. But it's one room in a very large house.

And while everyone's focused on that room, the rest of the house is quietly falling apart.

The contact centre problem

Contact centre KM is visible. It has vendors. It has metrics. It has a clear ROI story — faster handle times, fewer escalations, happier customers. There are software platforms built specifically for it. Conferences dedicated to it. Entire consulting practices around it.

So when organisations say "we have KM," they often mean "we have a tool that helps agents answer questions."

That's not KM strategy. That's knowledge delivery at one endpoint.

What's happening everywhere else

While the contact centre is meticulously managed, here's what's happening in the same organisation:

A senior leader retires and takes twenty years of institutional memory with them. Nobody captured it. Nobody thought to.

A project team spends three weeks solving a problem that another team solved eighteen months ago. The solution existed. It just wasn't findable.

A proposal goes out with assumptions that contradict positions taken on a previous engagement. The left hand didn't know what the right hand had written.

A new hire takes six months to reach full productivity because the knowledge they need lives in people's heads, not in any system.

Photo by Susan Q Yin on Unsplash

None of this shows up in a contact centre dashboard. All of it is a knowledge management failure.

The invisible KM problem

Contact centre KM is *visible* KM. It has a home, a budget, and a team accountable for it.

Everything else is *invisible* KM - the organisational knowledge that lives in people, relationships, past work, and institutional memory. It has no home. No owner. No metrics. And so, no urgency.

Until someone leaves. Until a bid is lost. Until a mistake is repeated. Until the organisation realises it keeps reinventing the wheel because nobody was minding the wheel.

KM is a whole-organisation discipline

The international standard for KM - ISO 30401 - doesn't mention contact centres. It talks about strategy, culture, learning, leadership, and governance. It defines KM as a management discipline that helps organisations create, retain, share, and apply knowledge to achieve their objectives.

APQC, one of the most respected research bodies in the field, frames KM across five pillars: strategy, process, content, culture, and technology. Technology is one pillar. The contact centre sits inside technology. That's a lot of building left unattended.

The questions that reveal the real problem

You don't need a formal audit to see where your organisation's invisible KM is failing. Just ask:

What happens when your most experienced person walks out the door?

How long does it take a new hire to become genuinely productive?

How do you capture what worked - and what didn't - at the end of a major project?

Can you find your organisation's best thinking on any given topic in under five minutes?

Do your teams build on each other's work, or do they start from scratch every time?

If those questions make people uncomfortable, that's not a contact centre problem. That's a knowledge strategy problem.

What to do about it

Start by separating the *tool* from the *discipline*. Contact centre software is a tool. KM is a discipline that applies across every part of an organisation - from how you onboard people to how you run retrospectives to how you capture competitive intelligence to how you make expert knowledge accessible to everyone who needs it.

Then look at where knowledge is genuinely at risk. Who are your critical knowledge holders? What happens when they leave? Where does important knowledge live that isn't written down anywhere?

You don't need a massive programme to start. You need someone asking the right questions. And you need leadership that understands that knowledge isn't just a contact centre asset — it's the organisation's most valuable and most under protected resource.

The contact centre KM is working fine. That's not the problem.

The problem is the illusion of coverage it creates. The belief that because one type of knowledge is well-managed, the rest is too. It isn't. And the cost of that gap shows up slowly - in turnover, in repeated mistakes, in lost bids, in the quiet erosion of things that once worked.

KM done well is invisible too. You notice it in how fast people find what they need. In how smoothly transitions happen. In how teams learn from each other rather than around each other.

That kind of KM doesn't live in a single tool. It lives in how an organisation thinks about what it knows - and what it can't afford to lose.

Stephanie Barnes is a Knowledge Management Consultant based in Canada who helps domestic and global organizations to increase customer satisfaction, and drive performance through re-engineering their knowledge management processes. Article originally published here.

Kapish expands with iCognition acquisition

Australian technology firm Kapish has acquired iCognition, bringing together two of the country's better-known names in information management and records governance.cKapish, which was founded in 2007 and operates as part of the Citadel Edge group, provides solutions across Information Management, Enterprise Architecture and Cyber Security to support customers through business transformation.

Founded over 20 years ago, iCognition is a leading provider of information management and governance solutions, becoming widely recognised as a trusted partner across all levels of government, whilst forging a longstanding reputation for expertise in OpenText Content Manager. iCognition helps organisations modernise their records management environments.

The deal is positioned as a capability expansion rather than a rationalisation, providing clients with access to a broader range of products and solutions and accelerated innovation and development.cThe combined companies will offer increased capacity to deliver, now backed by a larger team of specialists.

Kapish Managing Director Ryan Harris said the two organisations share a common focus on helping customers extract value from their information assets.

"Bringing iCognition into Kapish is a natural and exciting step forward for our business," Harris said.

"Both organisations share a strong commitment to helping our customers unlock the full value of their information and operate more effectively in complex digital environments.

"By combining our capabilities, technology, and deep expertise in information management and cloud technologies, we will accelerate product innovation and enhance the services and solutions we deliver to customers.

"We are delighted to welcome the iCognition team to Kapish and look forward to what we build together."

OpenText Asia Pacific Senior Vice President Stephen McNulty welcomed the acquisition, noting both firms had established reputations for customer-centric delivery.

"Together, they are well positioned to accelerate innovation, strengthen delivery capability, and support organisations in managing their information with greater confidence and efficiency," said McNulty.

"We welcome this next phase of growth and look forward to continuing our partnership with Kapish as they expand their capabilities and impact in the market."

For more information, visit kapish.com.au

Make records capture easy to drive adoption

By Demos Gougoulas

After years of working with organisations on records and information management, one lesson stands out clearly: if you want records captured, make capture easy. Most organisations already understand why records matter. Compliance, risk management and accountability are rarely in question. The real issue is effort. When record capture feels like extra work, it simply doesn't happen consistently.

The failure point is almost always the same. Organisations rely on policies and good intentions, then expect people to remember what needs to be captured, stop what they're doing to file information elsewhere, manually name and classify records, and make judgement calls under time pressure.

Even highly engaged teams struggle in this model - not because they don't care, but because friction always wins in a busy environment.

The biggest improvements occur when record keeping is no longer a separate task but part of the normal workflow. When capture happens directly from the systems people already use, is automated or near-invisible, requires minimal decision-making and applies rules consistently in the background, adoption improves rapidly. People don't resist record keeping - they resist unnecessary steps.

Well-designed technology removes that friction. It reduces effort and helps people do the right thing without disrupting their work. That said, technology isn't magic. It doesn't replace the obligation for sound record keeping practices - it supports them. Clear policies, accountability and good information governance remain essential. When capture becomes simple, attitudes change quickly. Once users trust that records are being captured as they're created or received, that metadata and classification are handled automatically, and that they're unlikely to "get it wrong," record keeping fades into the background. And that's the goal.

The best capture solutions are the ones users barely notice. If I had to distil this into one piece of advice: make capturing records easier than not capturing them. Not through more training, reminders or enforcement - but by implementing tools that respect how people actually work.

That's where I've consistently seen the biggest gains, and why ease of capture should always be the starting point, not an afterthought.

Demos Gougoulas is Director of Sales and Marketing at EzeScan.

Payments Surge Globally as Ransomware Groups Multiply

Ransom payments climbed sharply in 2025, with 24% of ransomware victims paying - up from 14% the previous year - as the number of active threat groups rose 16% to 67, according to a new global report.

The S-RM and FGS Global Cyber Incident Insights Report 2026, drawing on data from more than 800 incidents responded to globally in 2025, found the average ransom payment reached USD $296,000. Ransomware accounted for 45% of all incidents.

Asia-Pacific recorded the biggest regional surge, with over 760 organisations named on ransomware leak sites - a 59% increase on 2024. East and South-East Asia saw a 71% rise, the highest of any region globally.

US-based businesses remained the primary target, accounting for more than 60% of incidents. The report noted 45 unique threat actors targeted American companies - more than the rest of the world combined.

Despite the well-documented risk, basic security controls remain widely undeployed. Only 22% of ransomware victims had fully rolled out and actively monitored endpoint detection and response (EDR) tools across their environments.

VPN vulnerabilities continue to be the most exploited entry point. Single-factor remote access solutions accounted for 34% of ransomware entry methods in 2025, while public-facing infrastructure vulnerabilities accounted for a further 27.6%.

The Akira ransomware group alone was responsible for nearly 70% of Sonicwall-related incidents. VPN devices were the identified source in 68% of all remote access exploitation cases.

For Business Email Compromise (BEC) attacks - which made up 27.9% of all cases - credential phishing accounted for 80% of confirmed entry methods.

Of BEC victims, 47% had not enforced multi-factor authentication (MFA) in their Microsoft 365 environment. The average diverted funds from BEC attacks reached USD $165,000.

Double Extortion Now Standard Practice

Data exfiltration featured in 80% of ransomware attacks, as threat actors increasingly combine encryption with the threat of publishing stolen information to maximise leverage.

The report noted that data decryption is no longer the primary motivator for paying a ransom. In 88% of ransomware cases, victims had backups in place, with 69% having mostly viable backups - an improvement for the third consecutive year.

However, viable backups alone did not eliminate ransom payment: 30% of victims with fully viable backups still paid.

While 60% of victims chose to engage with threat actors - often to determine what data had been exfiltrated - only 41% of those who engaged ultimately paid.

Financial Services and Healthcare Most Exposed

Financial services topped the sector rankings, accounting for 12.7% of all incidents. Of financial services victims, 56% had no EDR deployed at all - nearly double the cross-sector average of 34%.

Healthcare was targeted by 21 unique threat actors - the highest of any sector - and also struggled with MFA deployment on remote access solutions.

Industrials and manufacturing recorded the highest ransom payment rate at 37%, compared with the 24% average across all sectors, reflecting the acute operational disruption ransomware causes in those environments.

Number of ransomware cases by country

in 2024 and 2025

Country	Ranking in 2025	Victims in 2025	Victims in 2024	Percentage Change
United States	1	4,057	2,783	46%
Canada	2	420	300	40%
Germany	3	312	169	85%
United Kingdom	4	278	266	5%
Italy	5	170	145	17%
France	6	168	125	34%
Spain	7	153	107	43%
Australia	**8**	**130**	**102**	**27%**
Brazil	9	127	121	5%
India	10	113	105	8%

Source: ecrime.ch

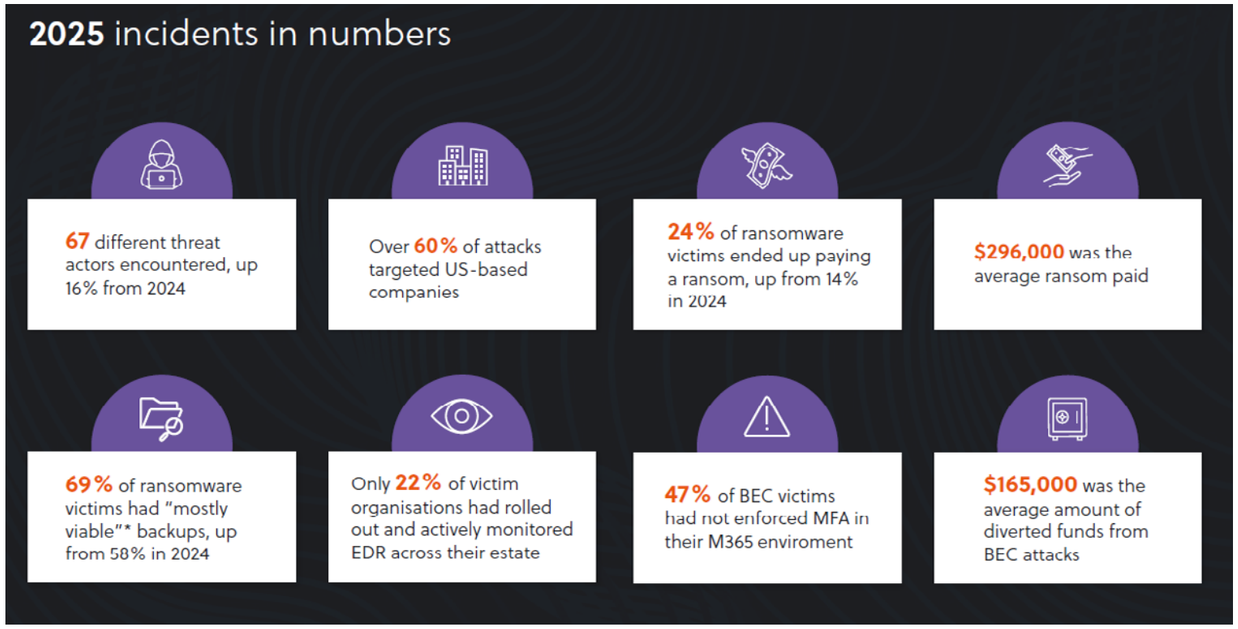

Professional services firms, particularly law firms, saw a higher-than-average rate of BEC attacks - 49% of cases versus a 28% average - due to frequent sharing of document links with clients.

Australia Under Escalating Pressure

Australia ranked eighth globally for ransomware victims in 2025, with attacks rising 27% year-on-year to 130 disclosed cases. Small and medium enterprises accounted for 78% of Australian organisations named on ransomware leak sites.

The Australian Cyber Security Centre reported that the cost of ransomware incidents for SMEs rose 14% to AUD $56,000. National carrier Qantas suffered a significant data breach affecting millions of customers during 2025.

Australia's Cyber Security Act introduced mandatory reporting requirements for ransom payments to the Australian Signals Directorate within 72 hours - the first legislation of its kind globally. The obligation applies to all organisations with an annual turnover of AUD $3 million or more operating in Australia, regardless of ownership.

Fragmented Groups, Unpredictable Outcomes

The ransomware ecosystem continued to fragment. Established ransomware-as-a-service operators Akira and Qilin together accounted for 45% of incidents. However, 61 other threat groups were also active, many operating with limited experience and unpredictable tactics.

The report identified new groups making Asia-Pacific organisations a strategic focus. Group NightSpire (101 total victims, 34% in Asia) and Dire Wolf (56 total victims, 50% in Asia) were among the most active new entrants.

"The landscape is also shifting: English-speaking threat actors claimed high-profile targets, AI-enhanced communications made established groups and lone operators more effective across borders, and attacks surged across Asia-Pacific," the report's authors wrote.

AI Agents Expanding Enterprise Attack Surface

The report identified insecure AI adoption within enterprises - not just AI-assisted attacks - as a primary emerging risk. As organisations deploy AI agents with access to email, files and integrated systems, they create new non-human identities with broad privileges.

AI agents are vulnerable to prompt injection attacks - where malicious instructions are embedded in data the agent processes. The report cited multiple real-world examples of agents being manipulated into exfiltrating data or facilitating account takeovers.

The report also noted that the report cited Anthropic's own disclosure of a case in which its Claude chatbot was allegedly used to carry out automated end-to-end ransomware attacks.

The authors concluded that fully autonomous attacks are not yet driving financially motivated incidents at scale, but warned that AI is compressing attack timelines and lowering the technical barrier for amateur threat actors.

A cautionary example was the January 2026 launch of OpenClaw, an open-source autonomous AI agent. The report described it as "riddled with vulnerabilities," with agents tricked into divulging passwords, API keys and downloading malware - yet it was installed hundreds of thousands of times.

"In 2026, organisations will do well to apply the same identity, privilege, and monitoring discipline to AI systems that they apply to human users," the report stated.

What to Expect in 2026

The report identified five key trends for 2026:

AI agent adoption will expand attack surfaces and complicate incident response, requiring new forensic frameworks for non-human identities.

Extortion will become more targeted, with AI-assisted triage of stolen data used to identify the most legally and reputationally damaging material.

Disrupted ransomware groups will rebrand rather than disappear, with operators retaining capabilities across law enforcement actions.

Ransomware execution speeds will accelerate further, compressing defender detection and response windows.

Unsecured VPNs will remain the most reliable entry vector until organisations fully adopt zero-trust architectures.

The full report is available at https://www.s-rminform.com/cyber-security/cyber-incident-insights-report.

3-Year Transition for AML/CTF Compliance

The Department of Home Affairs and AUSTRAC have announced transitional rules providing existing reporting entities a three-year period to transition customer due diligence obligations under Australia's reformed anti-money laundering regime.

The transitional rules allow existing reporting entities to continue using current customer identification procedures until 30 March 2029 or adopt reformed initial customer due diligence obligations at any point between 31 March 2026 and the transition deadline.

"The transitional rules will ensure the reforms work effectively in practice and will allow additional time for certain reporting entities to develop systems and processes and to meet certain new obligations," according to AUSTRAC's update published on its website.

The reforms commence 31 March 2026 following passage of the Anti-Money Laundering and Counter-Terrorism Financing Amendment Act 2024.

Existing reporting entities must implement ongoing customer due diligence obligations from 31 March 2026 with no transition period, according to the guidance.

The transitional period applies only to initial customer due diligence. During the three-year transition, reporting entities must fully comply with whichever method they select for all new customers and customer types.

McCullough Robertson partner Lucy Adamson notes the deadlines provide "implementation flexibility" while businesses must understand which timeframes apply to them, according to legal analysis published by the law firm.

The transitional period does not apply to newly regulated businesses that commence enrolment from 31 March 2026, including lawyers, accountants, conveyancers, real estate professionals, precious metals dealers, and trust and company service providers.

Existing reporting entities have until 30 May 2026 to notify AUSTRAC of their AML/CTF compliance officer. Newly regulated businesses and virtual asset service providers have until 29 July 2026.

Newly regulated businesses receive staggered deadlines for initial independent evaluations based on their AUSTRAC account number, avoiding sector-wide evaluation clustering.

"The deadline for the first independent evaluation for newly regulated businesses will not be less than 3 years from the original commencement date," the guidance states.

"This means the first deadline for a newly regulated business will be required to complete an independent evaluation will be 1 July 2029."

The reforms introduce regulation for new virtual asset services beyond the existing digital currency exchange framework established in 2018. The transitional rules defer obligations for these new services until 1 July 2026, aligning with broader tranche two reforms.

The Department of Home Affairs will publish an exposure draft of the transitional rules in coming weeks with industry feedback opportunities.

Reporting entities managing compliance processes face significant system development and workflow automation requirements under the reformed regime.

The extended timeframes provide breathing room for organisations implementing new data management systems and automated compliance workflows.

The reforms particularly affect information management systems requiring integration of ongoing customer due diligence monitoring, automated risk assessment frameworks, and enhanced reporting capabilities.

Organisations can direct questions about transitional rules to economiccrime@homeaffairs.gov.au.

2026 $7.5B Security Spend Signals AI Arms Race: Gartner

Deepening AI adoption and a widening talent shortage are combining to push Australian information security spending past AU$7.5 billion in 2026, a 9.5 per cent increase on 2025, according to new Gartner research.

Security software is the fastest growing segment, forecast to rise 12.3 per cent to more than AU$3.3 billion, driven by demand for application security, data security and privacy tools, and infrastructure protection.

Security services - encompassing consulting, professional services and managed security service providers (MSSPs) - remain the largest single spending category. Gartner forecasts the segment will reach more than AU$3.7 billion in 2026, up 6.9 per cent.

Analysts attribute this growth partly to organisations outsourcing security to fill capability gaps they cannot hire for domestically.

"This growth is fuelled by a growing, and increasingly critical, need for AI-literate security personnel, which is amplifying the continuing cybersecurity talent shortage in Australia," said Gartner VP Analyst Richard Addiscott.

"This is resulting in organisations needing to depend more on managed security service providers to fill skill gaps and enhance their security capabilities."

The talent shortfall is well documented. Independent research by ISACA found 54 per cent of Australian cybersecurity teams are understaffed, with 58 per cent reporting unfilled positions in 2025.

The Australian Computer Society's 2025 Digital Pulse report estimates 54,000 more skilled cyber professionals will be needed by 2030.

CyberCX and Per Capita research has forecast a shortfall of up to 30,000 unfilled cyber positions across Australia in coming years.

Gartner predicts more than 75 per cent of enterprises will be using AI-amplified cybersecurity products across most use cases by 2028, up from less than 25 per cent in 2025.

The firm says vendors are embedding AI-driven threat detection and automated incident response into mainstream product lines.

Addiscott described GenAI tools as increasingly deployed because security teams "leveraging traditional security measures struggle to scale and keep pace with a constantly evolving threat environment."

https://www.gartner.com

Rethinking the Essential Eight: Cybersecurity in the Age of AI

By Ghaith Kayed

As someone working at the intersection of cybersecurity and public sector technology, I've long respected the Essential Eight framework developed by the Australian Cyber Security Centre (ACSC). It's practical, actionable, and has helped lift the security posture across government agencies and critical infrastructure. But the world has changed. And so must our approach.

In this short article, I'm hoping to start a conversation and offer some practical ideas for how we can evolve the framework in a way that keeps pace with AI-driven threats while preserving its core strengths.

This article is not about discarding what works. But about building on it. The Essential Eight has been a cornerstone of cyber hygiene in Australia. But in a world of AI-powered threats and AI-dependent systems, we need to ask: **Is it enough? And how do we evolve it while keeping it practical and widely adoptable?**

The core message: Why the Essential Eight needs to evolve

The Essential Eight was designed for a threat environment dominated by conventional malware, phishing, and privilege escalation. Today, attackers are using AI to:

- Rapidly generate polymorphic malware.
- Craft highly convincing phishing at scale.
- Bypass traditional application controls.
- Target AI models and the data that feeds them.

Meanwhile, governments are increasingly considering using AI to support making decisions, managing infrastructure, and delivering public services making those systems targets in their own right.

If we're going to defend in this new era, we need to update the playbook.

The new AI landscape

AI is not just reshaping how we work. It's reshaping how attackers operate. We are all seeing:

- AI-generated malwares that can mutate faster than signature-based tools ability to catch.
- Social engineering campaigns scaled by generative language models.
- Deepfakes that mimic trusted identities.
- New attack surfaces across machine learning models and data pipelines.

The Essential Eight was not built for this reality. I'm sharing this to spark a conversation:

How do we evolve the frameworks we trust without losing their simplicity or clarity?

I don't have all the answers, but I believe this is the right time to ask better questions.

- How do we modernise our most trusted frameworks without overcomplicating them?
- What is already working that we can learn from?
- What risks and opportunities do **you** see in applying AI to both defence and offense?

If you work in or around cyber strategy, government systems, or critical infrastructure, I'd love to hear your input. There's an opportunity to collectively evolve the frameworks that keep our systems safe.

1. A Proposal: Expanding to an "Essential Ten"

To meet the challenges of the AI era, I believe we need to expand the framework to include two new strategies. They're extensions of core security principles adapted to a new class of assets: **AI models and training data**:

- **AI system integrity:** As government agencies deploy AI to support decision-making, fraud detection, or service delivery, we must secure the models themselves. That means testing for adversarial inputs, monitoring for drift, securing model pipelines, and validating training data.
- **Data provenance and lineage:** AI systems are only as trustworthy as the data they learn from. Without proper tagging, lineage tracking, and origin checks, we risk training sensitive systems on poisoned, biased, or unauthorized data.

These aren't just IT hygiene issues. They're national security concerns.

2. Enhancing the original Eight to be AI-Aware

Every one of the existing Essential Eight strategies can (and should) be updated to account for AI-enhanced threats. Keeping the "essential" truly essential but making it current.

For example:

- **Application control** must detect evasive, AI-generated binaries.
- **Multi-factor authentication** needs to go beyond passwords + SMS to include continuous authentication, behavioural biometrics, and phishing-resistant tokens.
- **Backup strategies** must include validation to ensure AI-corrupted data hasn't silently made its way into recovery points.
- **Patch management** should leverage AI-driven threat forecasting. Not just CVSS scores.

3. Move from static compliance to continuous assurance

In an AI-driven threat environment, annual audits or static checklists aren't enough. We need **realtime, automated, AI-assisted validation** of security controls. This allows us to detect gaps before they become breaches and better allow responding in hours, not weeks.

4. Continue supporting a shared AI threat intelligence fabric

Government and critical infrastructure sectors can easily become siloed. We need to continue supporting a secure, cross-agency intelligence sharing network that can use AI to correlate signals and early identify threats without compromising data privacy.

This might include federated learning approaches, realtime telemetry sharing, or red/blue teaming with AI agents.

5. Make AI security a core part of cyber culture

We must build **AI security literacy into the culture** of cybersecurity teams. Defending against AI-powered attacks and protecting AI systems requires new knowledge and skills. This includes understanding adversarial ML, data poisoning, model inversion, and much more.

We need to train for the world we're entering, not the one we're leaving behind.

6. Make the Essential Eight framework collaborative and open (Within trusted bounds)

Today, the Essential Eight is centrally managed. While that ensures consistency, it can also limit agility.

What if we opened it up (at least domestically) for contribution by trusted experts across government, industry, and academia?

A secure, transparent model for collaborative evolution similar to opensource softwa.re but with tiered review and approval. This could help:

- Tap into the collective intelligence of the cybersecurity community.
- Respond to emerging threats faster.
- Build shared ownership over a framework that protects us all.

In the age of AI, frameworks can't remain static. We need mechanisms to evolve in realtime.

Ghaith Kayed has 20 years of experience delivering AI, IoT, and analytics programs across Australia, New Zealand, UK and the USA. Article originally published here.

	Strategy	AI-Age Enhancement
1	Application control	Detect AI-generated evasion techniques
2	Patch applications	Add AI-powered threat forecasting
3	Office macro settings	Detect AI-assisted phishing/doc crafting
4	User app hardening	Defend against AI-manipulated browser exploits
5	Admin privilege restrictions	Detect AI-driven privilege escalation attempts
6	Patch operating systems	Automate with AI-based exploit detection and prioritization
7	Multi-factor authentication	Upgrade to continuous and contextual authentication
8	Regular backups	Ensure AI-generated/modified data is excluded from clean backups
9	**AI System integrity**	Secure AI model lifecycle and defend against adversarial inputs
10	**Data provenance and lineage**	Ensure AI training data is traceable, clean, and policy-compliant

Example: Enhanced Essential Eight evolving in the age of AI

New Zealand Proposes Mandatory Cyber Security Regime

The New Zealand Government has released a discussion document proposing mandatory cyber security obligations for operators of critical infrastructure, including enforceable minimum standards, incident reporting requirements, and director-level accountability.

Published in February 2026 by the Department of the Prime Minister and Cabinet (DPMC), the document warns that cyber risks "are generally not well understood or collectively managed to a consistent level across New Zealand's critical infrastructure system" (Prime Minister Christopher Luxon).

nternational concern about state-sponsored cyber threats. In October 2025, New Zealand's National Cyber Security Centre (NCSC) joined international counterparts in alerting organisations to Salt Typhoon - a People's Republic of China (PRC)-affiliated threat group. The document states that "Salt Typhoon activity has been observed in New Zealand", targeting critical infrastructure for espionage and potential sabotage.

New Zealand currently ranks 49th on the National Cyber Security Index - the lowest of all Five Eyes partners - and sits in the third tier of the Global Cybersecurity Index, while partner nations occupy the first tier.

The document states New Zealand "stands out from other advanced economies in not using dedicated legislative mechanisms to protect critical infrastructure from cyber harm."

Scope of the Proposed Regime

The proposed regime would apply to approximately 200 entities across seven essential service sectors: communications and data, defence, energy, finance, health, transport, and drinking water and wastewater.

Thresholds defining which entities are captured include electricity generators with capacity at or above 30 megawatts, registered banks identified as domestically systemically important, hospitals with intensive care units, telecommunications operators serving at least 10,000 customers, and maritime ports handling more than 4 million tonnes of freight annually.

A smaller subset - designated as critical infrastructure of national significance (CINS) - would face additional obligations. Designations would be made privately, not disclosed publicly, for security reasons.

Six Proposed Measures

The document proposes six measures, which can be adopted individually or as a package.

- **Measure 1** would grant the responsible Minister power to require critical infrastructure entities to provide operational information to government, including details of critical components, ownership and control structures, and mapping of dependencies.
- **Measure 2** would establish a voluntary cross-sector information exchange to connect critical infrastructure entities with each other and with government, enabling coordinated cyber security responses.
- **Measure 3** would require entities - initially those designated as critical infrastructure of national significance - to share specified information with each other, such as projected restoration times.
- **Measure 4** would introduce mandatory cyber incident reporting to the NCSC, including an initial early warning within 24 hours and a full report within 72 hours for significant incidents. Significant incidents are defined as those having, or likely to have, serious impact on the confidentiality, integrity or availability of information, or on the delivery of essential services.
- **Measure 5** - the centrepiece for compliance managers - would require entities to develop, implement and maintain a risk management programme aligned with an internationally recognised cyber security framework, such as the NIST Cybersecurity Framework or ISO/IEC 27001:2022. The programme must identify critical components, assess material risks, and treat those risks as far as reasonably practicable.
- **Measure 6** would grant the Minister a last-resort power to direct a critical infrastructure entity to take - or refrain from - specified actions to manage a cyber threat for national security reasons. The document states this power would only be used where "the national security threat is significant" and "there is no alternative to the proposed action that would satisfactorily address the national security threat."

Director Accountability and Penalties

Directors of critical infrastructure entities would bear personal responsibility for compliance with minimum requirements. The document proposes making cyber security a core element of directors' fiduciary duty.

Penalties for the most serious breaches would include criminal fines of up to $NZ5 million, or 2 per cent of annual turnover (whichever is greater), for entities, and up to $NZ500,000 for individual directors. Minor breaches may attract administrative fines starting at $NZ50,000. A staged compliance approach is proposed, with a one-year grace period before enforcement action is considered. Third-party audits are considered unlikely in the medium term, due to cost and limited market capacity.

Supply Chain and Third-Party Obligations

The proposals extend obligations to suppliers and contractors that have operational control over critical components. Third-party vendors would be required to support critical infrastructure entities in meeting their risk management obligations, as far as reasonably practicable.

This has significant implications for managed service providers, cloud computing vendors, and data centre operators serving the critical infrastructure sectors. The document proposes that data centres and managed service providers integral to the delivery of essential services by a critical infrastructure entity would themselves be designated as critical infrastructure.

The supply chain risk was highlighted by the 2024 CrowdStrike software update incident, which caused widespread IT outages across finance, healthcare and transport sectors globally, and the Manage My Health breach that compromised personal data of up to 126,000 New Zealanders. https://www.ncsc.govt.nz/news/rise-in-financial-losses-reported-to-the-ncsc

Claude Mythos and Project Glasswing

Why an AI superhacker has the tech world on alert

Project Glasswing, named for the Greta oto butterfly, Westend61 / Getty Images

By Stan Karanasios and Saeed Akhlaghpour, The University of Queensland

New, more powerful artificial intelligence (AI) models are announced pretty regularly these days: the latest version of ChatGPT or Claude or Gemini always has new features and new capabilities that its makers are eager for customers to try out.

But now Anthropic has announced a new model with great fanfare, but is only giving access to a select handful of users. In what the New York Times calls a "terrifying warning sign" of the model's power, the company has instead started an initiative called Project Glasswing to use the model for good instead of evil.

Why? Early reports indicated that the model, with instruction, had been able to move outside a contained testing "sandbox" and send an email to a researcher.

A little alarming, perhaps. But more significantly, Anthropic claims Mythos has uncovered software vulnerabilities and bugs "in every major operating system and every major web browser".

Finding hidden vulnerabilities

In one remarkable example, the model found a flaw in OpenBSD, a security-focused operating system used in firewalls and routers, which had gone undetected for 27 years. According to Anthropic, it also found a 16-year-old vulnerability in FFmpeg, a little-known but widely used behind-the-scenes piece of software that helps computers, apps, and websites handle audio and video files.

Anthropic also says Mythos found several vulnerabilities in the kernel of the Linux operating system, and chained them together in a way that could give an attacker complete control of a machine.

Anthropic's internal assessment of the model highlights both its technical promise and the need for vigilance.

The report outlines a hypothetical risk that an advanced AI might exploit its access within an organisation, but concludes that the model poses a very low threat of harmful autonomous actions. In other words, it is unlikely to "go rogue" – but may follow human directions to do things that cause harm.

Why Anthropic is keeping Mythos off-limits

Anthropic says it decided not to release the model publicly because of its capabilities and the potential risks it poses. At the same time, the company launched Project Glasswing.

The effort brings together a broad coalition of tech companies such as Microsoft, Amazon, Google, Apple, Cisco and NVIDIA, open-source organisations such as the Linux Foundation, and major financial actors such as JPMorganChase, to channel Mythos towards cyber defence rather than misuse.

The idea is to give defenders a head start to find and fix weaknesses in critical software before similar AI capabilities become widely available to attackers.

Reading between the lines of Anthropic's messages

This is not the first time an AI firm has decided a model was too powerful to release widely. In 2019, years before the ChatGPT era, OpenAI did something similar with its (now quite primitive-looking) GPT-2 model. (Dario Amodei, now chief executive of Anthropic, was a key OpenAI researcher at the time.)

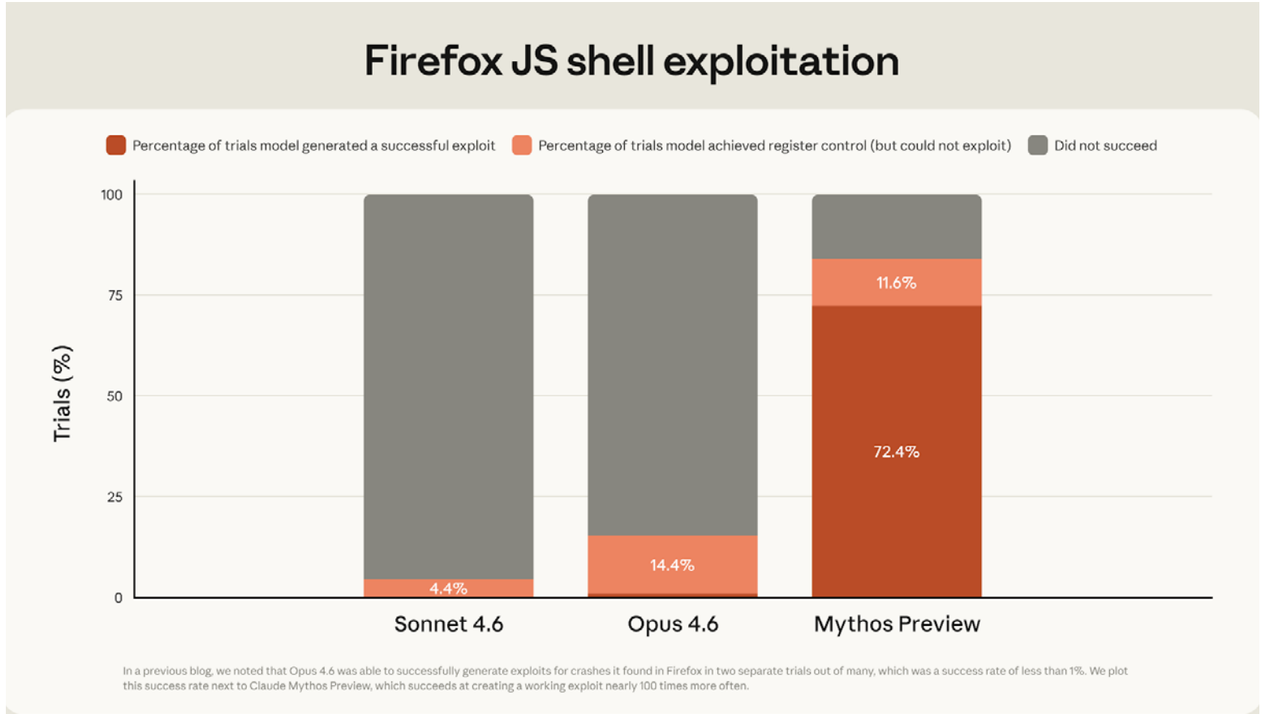

Anthropic's internal testing (which has not been independently verified) showed the Mythos model was far more successful than earlier models at turning software bugs into working exploits. Anthropic

However, this doesn't mean these announcements should not be taken seriously.

Anthropic has published unusually detailed material for a model it is not widely releasing. Reports suggest US authorities convened major US bank CEOs in Washington to discuss the cyber risks associated with Mythos.

However, we should exercise caution about Anthropic's claims, because outsiders cannot yet verify most of the underlying evidence. Anthropic says more than 99% of the vulnerabilities it found are still undisclosed because they have not yet been patched. That is responsible disclosure, but it also means the public is being asked to trust a great deal it cannot fully inspect.

What Mythos could mean for the future of cybersecurity

Cybersecurity failures can have real effects on individuals. In Australia, the Optus breach exposed the personal information of about 9.5 million people. In another case, stolen Medibank records included sensitive health information, and some of the data was later released on the dark web.

These were not just database problems. They became crises of privacy, identity and trust.

That is why Mythos matters. Mythos and other AI models like it could change the basic economics of cybersecurity.

In the past, serious vulnerabilities have often stayed hidden simply because nobody found them. And this in turn was because finding them took rare skill, patience, and time.

If models like Mythos can scan the hidden plumbing of the internet – operating systems, browsers, routers, and shared open-source code – at an unprecedented scale, then what is now specialised hacking could become a routine and automated process.

For organisations and software development firms, Mythos is a double-edged sword. It could rapidly uncover hidden flaws in their own code, but it also raises the fear attackers could find the vulnerabilities first.

The implications reach well beyond tech companies. Much of that underlying, invisible software supports many of the services people rely on every day, from electricity and water to airlines, banking, retail and hospitals.

What now?

So far, cybersecurity and software companies have been remarkably quiet in public about Anthropic's Mythos. Many firms appear to be waiting and watching, unwilling to signal their stance in case the model exposes weaknesses in their own systems.

But developments like Mythos are a reason to stop treating cybersecurity as somebody else's problem. For now, for individuals, the response is simple: basic cyber hygiene matters more than ever.

Update phones, laptops, browsers and routers. Replace unsupported devices. Use a password manager. Turn on multi-factor authentication. Do not ignore patch notices.

Those are the immediate steps. Beyond them lies a harder set of questions about AI and cyber security – about who gets access to powerful AI models, who oversees their use, and who decides what counts as the "right hands".

Stan Karanasios, is Professor in Information Systems, and Saeed Akhlaghpour, is Associate Professor of Business Information Systems, The University of Queensland.

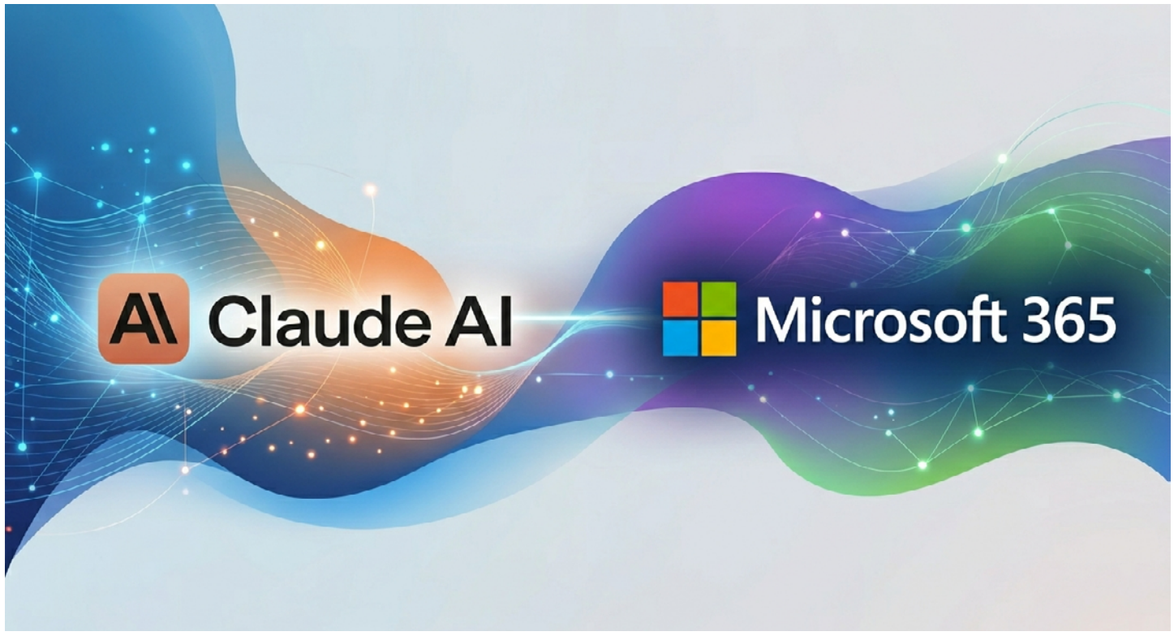

Claude is now inside your M365 tenant, but mind the data residency gap

Copilot Cowork is Microsoft's new cloud-based AI agent built with Anthropic's Claude technology that can plan, execute, and deliver multi-step work across Outlook, Teams, Excel, PowerPoint, and SharePoint on a user's behalf.

The product is the centrepiece of what Microsoft is calling Wave 3 of Microsoft 365 Copilot — a sweeping platform update that also brings Anthropic's Claude models into mainline Copilot Chat and introduces new enterprise pricing tiers. It draws directly on the same agentic harness as Anthropic's Claude Cowork, which had already been released for Mac in January 2026 and Windows in February.

Unlike Anthropic's desktop version, which runs locally on a user's device, Copilot Cowork operates entirely in the cloud within a customer's Microsoft 365 tenant. This means it is covered by Microsoft's enterprise data protection framework and connected to what Microsoft calls Work IQ — an intelligence layer that draws on the full graph of a user's work data, including email threads, Teams conversations, calendar history, SharePoint files, and Excel workbooks.

Microsoft president of Business Applications and Agents, Charles Lamanna, described the shift in a company blog post: "Copilot Cowork is built for that: it helps Copilot take action, not just chat." Users describe the outcome they want, and Cowork breaks the request into a structured plan that runs in the background, surfacing checkpoints for approval before any changes are applied.

In a demonstration, Lamanna showed Cowork analysing a month of meetings with direct reports, compiling customer notes from a business trip, and generating a competitive analysis with an accompanying Word document and Excel spreadsheet — without requiring manual coordination across apps.

The Anthropic partnership

The launch reflects the deepening relationship between Microsoft and Anthropic, formalised through a US$30 billion Azure compute deal announced in November 2025. Anthropic has operated as a Microsoft sub-processor since January 2026, covered by Microsoft's Product Terms, Data Processing Agreement, and Enterprise Data Protection framework.

Microsoft chief marketing officer for AI at Work, Jared Spataro, said the multi-model approach was a deliberate differentiator: "Every 60 days at least, there's a new king of the hill. There's so much demand for a platform that doesn't feel like, 'I have to skip over to the next vendor."

Claude handles the complex reasoning components of multi-step tasks, while Microsoft's own models manage integration with M365 applications.

Fortune reported that Spataro drew a direct distinction between Anthropic's consumer product and the enterprise offering, describing the desktop version as a "fantastic tool" but one with "limitations" in a corporate environment, given its lack of access to cloud-based enterprise data and governance controls.

Availability and pricing

Copilot Cowork entered Research Preview on 9 March with a limited set of customers. Broader availability through Microsoft's Frontier program was expected in late March 2026. Accessing the feature requires a Microsoft 365 Copilot licence at US$30 per user per month on top of existing enterprise subscriptions.

Microsoft also announced the Microsoft 365 E7 suite, which will be generally available from 1 May 2026 at US$99 per user per month. E7 bundles Microsoft 365 Copilot, Agent 365, the Microsoft Entra Suite, and Microsoft 365 E5 security capabilities under a single plan — the first new enterprise licence tier in approximately 10 years.

Agent 365, the governance and orchestration platform for AI agents, will also be available as a standalone add-on at US$15 per user per month from 1 May.

Claude models are currently available in Copilot Chat, Researcher, and Excel for tenants holding a Copilot licence, though administrators must enable the Anthropic model in Copilot settings. Microsoft notes that Copilot Cowork itself remains in Research Preview, and no firm commercial availability date for general rollout has been committed.

A specific gap for A/NZ users

For Australian and New Zealand organisations, the governance picture requires careful mapping before adoption.

The Anthropic sub-processor toggle was enabled by default for most commercial tenants on 7 January 2026. EU, EFTA, and UK tenants have it disabled by default. Government cloud environments have no access to the feature.

Microsoft made in-country Copilot processing available for Australia in late 2025, meaning standard Copilot interactions can be processed and stored within Australian borders.

However, Anthropic models are explicitly excluded from those in-country commitments. Any request routed through Claude — including Cowork task execution — falls outside that data residency guarantee.

For New Zealand organisations, the same commercial tenant rules apply: the sub-processor toggle is on by default, and there are no equivalent in-country processing commitments that would cover Anthropic model use.

John Hoy, Founder of Orbital Intelligence and a member of the Governance Institute of Australia, has posted a detailed analysis of the governance implications for Australian and New Zealand enterprises on LinkedIn.

Hoy identifies four questions that M365 Copilot organisations need to address before enabling the feature: whether the Anthropic sub-processor toggle has been reviewed; which Copilot features route to Claude; whether data flows have been mapped; and whether existing AI governance frameworks account for a multi-model environment.

"Your standard Copilot interactions can stay onshore, but the moment a request routes to Claude, that guarantee falls away," Hoy writes.

"If in-country processing underpins your compliance posture, this needs attention before adoption."

Hoy tested Copilot Cowork extensively in advance of the announcement and described the capability as a genuine productivity force multiplier, noting a file management task that would have taken three hours was completed by the agent in 47 seconds.

But he was unequivocal that capability should not pre-empt governance: "Impressive capability sure, but there's a real governance gap here worth getting right before adoption."

Governance and enterprise readiness

Independent analysis from enterprise consultancy ARO echoes those concerns. In a post-announcement commentary, Paul Ryder of ARO argued that while the capability is a significant step toward integrating agents into daily processes, most organisations are not structurally or operationally ready for this level of automation. "The readiness gap matters far more than the technology itself," he wrote, pointing to the need for data

"Your standard Copilot interactions can stay onshore, but the moment a request routes to Claude, that guarantee falls away.

"If in-country processing underpins your compliance posture, this needs attention before adoption."

quality, clear automation boundaries, defined ownership models, and training investment before broad adoption.

Governance researcher and enterprise security firm Pragatix identified an additional technical limitation: Copilot Cowork executes actions using the identity of the user rather than a dedicated AI agent identity, creating ambiguity in audit trails.

When an agent acts autonomously on a user's behalf — sending an email, modifying SharePoint permissions, updating a project timeline — existing compliance frameworks that assume a direct link between a user identity and a system action may not adequately capture accountability.

Microsoft has addressed the broader governance architecture in its product documentation. Cowork runs in a sandboxed cloud environment within Microsoft 365's security and compliance boundaries, with identity, permissions, and compliance policies enforced by default. Actions and outputs are auditable. The agent cannot perform actions that fall outside the permissions already held by the delegating user's account.

For those evaluating adoption, the immediate priority is understanding which Copilot features within their tenant currently route to Claude models, confirming whether the Anthropic sub-processor toggle has been reviewed against their data governance framework, and determining whether their AI governance policies have been updated to address multi-model agentic environments.

Market context

Microsoft's launch follows significant market disruption caused by Anthropic's standalone Claude Cowork releases in early 2026, which triggered a combined US$285 billion selloff in enterprise software stocks as investors repriced companies whose core workflows overlapped with what the desktop AI agent could automate.

With paid Copilot adoption still representing a small fraction of Microsoft's 450 million commercial Microsoft 365 subscribers, Cowork is widely seen as Microsoft's most significant effort to justify the Copilot licence cost and accelerate enterprise take-up.

Analyst commentary collected by Winbuzzer described the launch as Microsoft absorbing a competitive threat, repackaging Anthropic's technology within Microsoft's enterprise governance architecture, and using the resulting product to address both AI adoption and licence conversion.

Microsoft has also confirmed that Anthropic's Claude models are now available across the full Copilot Chat experience for Frontier program users, not solely in the Researcher and Excel features where Claude was previously integrated — meaning Anthropic's models are becoming a general-purpose option for everyday Copilot interactions, not only for task execution in Cowork.

Do you Believe in the Data Fairies?

By Nicola Askham

Picture the scene: you arrive at work, open your laptop, and settle down to do some important work, confident that the data you need will be ready and waiting for you. But here's my question: do you know how that data gets there, and where it comes from?

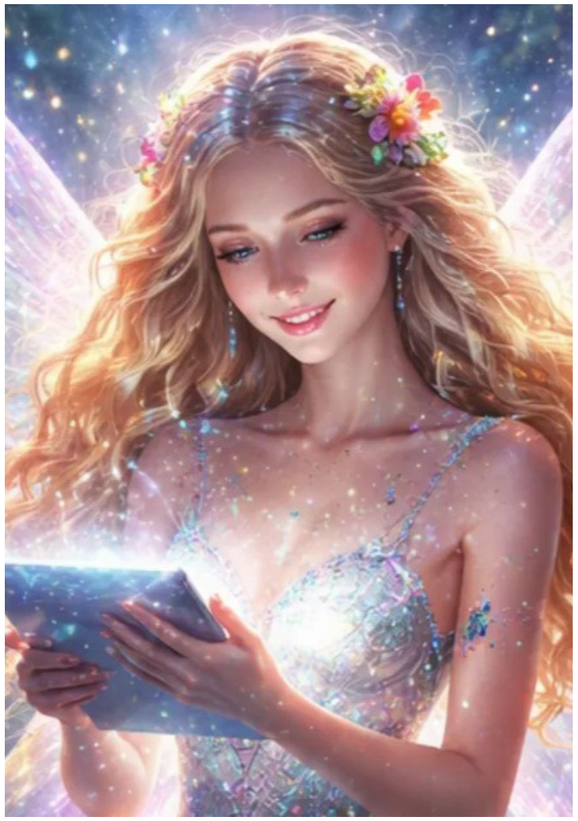

Over the years I've come across countless situations that run something like this:

Team A: Our data is loaded up by IT.

IT: No, we don't touch that data — it's a manual load by Team B.

Team B: We just send the spreadsheet to Team A. We're sure they load it.

Team A: No, we really don't load that data…

My friend and colleague Justin York coined the phrase "data fairies" for precisely these situations, and it describes them perfectly. For many people, the data simply *appears* and whilst they may not literally believe in the data fairies, they have no idea how they get their data, or sometimes even where it comes from. Sadly, they often don't care either.

I was lucky enough to hear Peter Aiken speak years ago, and he summed the situation up with a brilliant analogy: most people think about data the same way they think about air.

They don't think about it at all. They just assume it will always be there, and always be good enough to use. It's only when something goes wrong that they stop to consider the quality of what they're working with.

And I'd add: even then, do they think about *where* the data comes from?

All too often, the answer is no. Having been forced to acknowledge a data quality problem, the usual reaction is a tactical fix of the data in front of you. After all, if you don't know where the data came from or how it got there, how could you possibly consider a more permanent solution?

And so begins a cycle of constant data cleanses and patches that quietly become part of the normal way of working — nobody questions them, they just become "the process."

This is why **data lineage** (knowing where your data comes from and what happens to it on its journey to you) matters so much. It allows problems to be traced back to their source, and permanent solutions to be properly considered.

Now here's where 2026 changes everything.

When I wrote this post in 2012, the stakes of not knowing your data lineage were significant but manageable. Poor decisions, regulatory headaches, wasted hours fixing bad reports. Not ideal, but survivable.

In the age of Gen AI, the stakes are categorically higher.

Organisations are now feeding data into LLMs to drive automated decisions, about customers, about risk, about operations. If you don't know where that data comes from, you cannot know whether your AI is trustworthy.

Garbage in, garbage out has always been true, but when the "out" is an automated lending decision, a medical recommendation, or a fraud flag, the consequences are very different from a slightly wrong internal report.

There's a cruel irony here too. Many organisations are turning to AI to *help* with data quality and data lineage. It can genuinely provide some help for these activities, but AI cannot magic good data governance into existence.

It can help you document lineage faster, flag anomalies more quickly, and surface patterns a human might miss. What it cannot do is substitute for knowing, at a fundamental level, what your data is and who is responsible for it.

The data fairies are not just still with us, they're now feeding the models.

My advice remains the same as it was in 2012, but it's more urgent: take an iterative approach to building up your data lineage.

Use every data quality incident as an opportunity to document what you learn. Build a repository, piece by piece. And critically, make sure your AI governance and your data governance are not two separate conversations - because they cannot be.

The data fairies aren't going away on their own. In the age of Gen AI, understanding your data has never mattered more. My book Effective Data Governance (Kogan Page, 2026) gives you the framework to tackle all of this. It's available now.

Originally published here.

When Data Moves but Meaning Disappears - how AS5393 Can Help

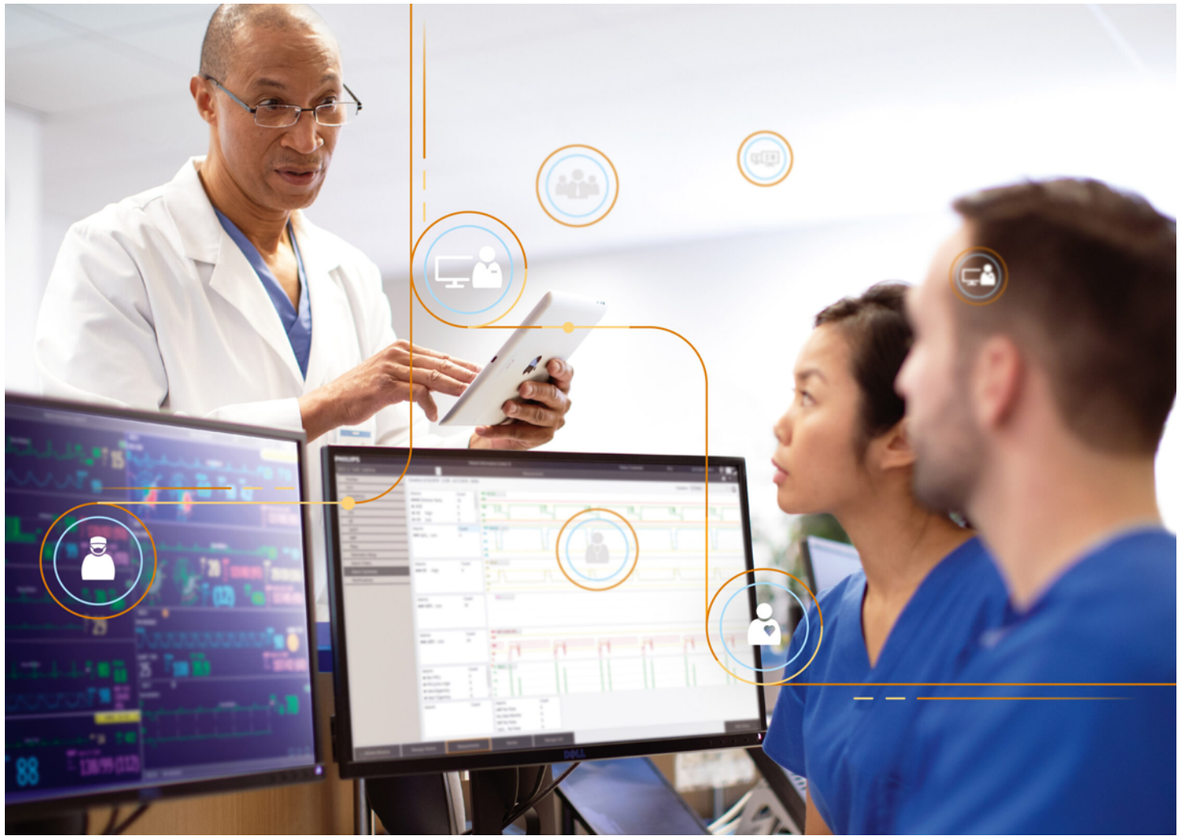

By Peta Sweeney, RIMPA Global

When a hospital migrated patient records to a new clinical management platform, the technical transfer succeeded. Every patient had a record in the new system. But six months later, clinicians discovered that specialist treatment notes were no longer linked to the diagnostic imaging that informed them.

The clinical context - the 'why' behind treatment decisions - had been broken. The health service spent eighteen months manually reconstructing relationships while managing emerging clinical safety risks.

In banking, a different migration failure unfolded. During an upgrade of a core banking system, transaction records migrated cleanly - balances, histories and timestamps intact. What did not survive were the metadata markers indicating legal holds.

The records still existed but the controls that restricted access to certain funds had disappeared, exposing the organisation to regulatory and anti-money-laundering risk.

These failures share a common cause. Records migration was treated as a technical exercise rather than as a governance and evidentiary problem. Australia's new AS5393 standard (*Records and information management - Migration of authoritative data, information and records between systems*) responds directly to this gap, reframing migration as a question of authority, context and trust rather than data movement alone.

What distinguishes AS5393 from existing records and information management standards is not that it introduces new principles but that it applies them explicitly to migration.

Rather than treating migration as an implementation detail to be managed by technology teams, the standard frames it as a discrete, high-risk activity that requires its own governance, planning, validation and assurance.

It does not replace established records management, information governance or system standards; instead, it operates across them, focusing on the point where policy intent, record-keeping requirements and system behaviour collide.

In doing so, AS5393 makes explicit the decisions that are often made implicitly during system change and requires those decisions to be documented, tested and defensible.

Records migration is not uniform. The consequences of failure differ distinctly across sectors, even when the technical task appears similar.

In government, migration is often driven by machinery-of-government change, system consolidation or long-term preservation requirements.

When a function transfers between agencies, records must migrate with it - not merely as files, but as evidence of lawful decision-making. Authority, provenance and administrative context must remain intact so that decisions made in one agency remain defensible in another.

In financial services, regulatory survivability dominates. Transaction records, advice documentation and compliance artefacts must maintain evidential integrity across system change.

A migration that strips metadata indicating legal holds or advice lineage may appear technically successful while creating immediate regulatory exposure.

In health, migration risk extends beyond compliance into clinical safety. Patient records are not static artefacts; they inform future care.

A pathology result without the clinical context of why it was ordered is not just incomplete - it is potentially dangerous. Migration decisions therefore become clinical governance decisions.

Legal practice brings different sensitivities. Matter files embed privilege, chain of custody and professional responsibility. A migration that intermingles privileged and non-privileged records or breaks internal matter structure, creates risk that no technical remediation can fully undo.

Architecture and engineering organisations operate on long liability horizons. Design calculations, approvals and as-built documentation form interdependent record sets that may need to be relied on decades later.

Losing relationships between drawings and calculations does not merely degrade records - it creates professional and legal risk long after project completion.

At the heart of AS5393 is a simple question: what makes a record authoritative once it leaves the system in which it was created?

The standard distils this into four principles that every migration must meet:

- avoid loss
- appraise and sentence
- ensure provenance
- maintain links.

While these principles apply universally, their meaning shifts by sector.

The highest-risk migrations are rarely simple upgrades. They occur during integration and decommissioning - precisely the scenarios most organisations now face as they consolidate platforms, move to cloud services or exit unsupported systems.

AS5393 does not prescribe a single migration method. It provides a framework for judgement. Used well, it enables organisations to preserve not just records, but the authority those records must continue to carry when systems - inevitably - change.

About the Standard

AS5393:2025 was developed by Standards Australia Committee IT-021 Working Group 15, with representation from government archives, universities, professional bodies and private sector organisations across Australia. It is the first international standard to specifically address records migration methodology and technical specifications, providing a structured framework for managing the risks associated with migrating authoritative data, information and records between systems.

Link to Standards Australia: https://store.standards.org.au/product/as-5393-2025

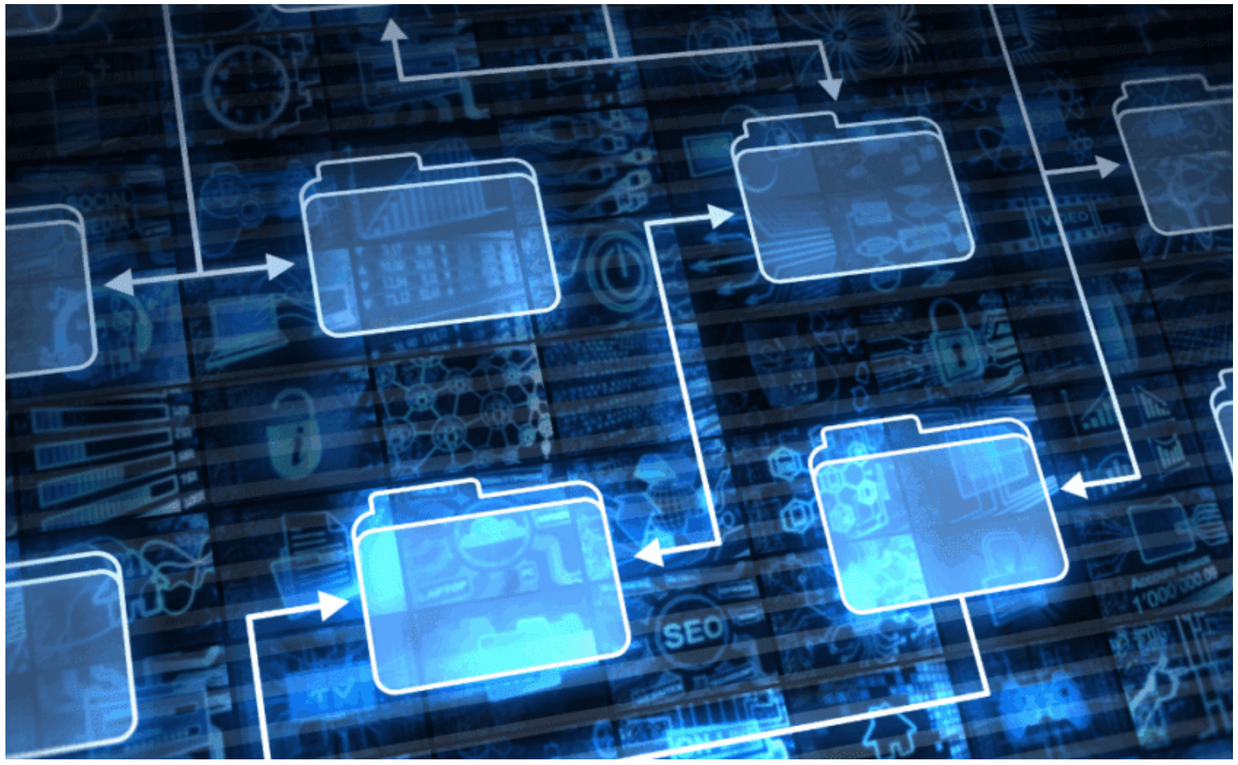

Many migration failures arise not from negligence but from system design assumptions. Platforms optimise for storage, retrieval or performance, while organisations rely on relationships, context and narrative continuity.

The Minister Tweeted. The CIO Didn't Sleep

By Newgen Software

The clock on his screen flickered to 3:16 AM. Rony Mehta, Chief Information Officer, for the state's citizen services department, rubbed his eyes and stared at the dashboard one more time. The new AI chatbot, launched just six weeks ago after a frantic development sprint, was processing permit applications 100x faster than his human team. The metrics were beautiful. The media loved it. The minister had tweeted about it.

So why couldn't he sleep?

Rony clicked open a random transaction. A permit application from a small business owner in a distant district. Approved in 39 seconds. He scrolled down, looking for answers. Where was the data processed? Mumbai? Singapore? Ghana? The audit log was... vague. How was this decision made? Just three words - Eligibility criteria met. That was it. No reasoning, no trail, no transparency.

Roni shut his laptop at 4 AM. He didn't have answers. But he knew, with cold certainty, that someone would soon ask the question. He just hoped he'd have an answer before they did.

The Problem wasn't his Alone

Across the world in 2024, governments raced to digitize. They launched intelligent chatbots, modernized forms, and automated workflows. These pilots proved technology could scale fast. The mandate was to **Get** AI **and get it real fast.** And, they did. They built Ferraris and drove them on goat trails, powerful systems on weak foundations.

However, as they moved into core public services, permits, pensions, healthcare, benefits, one realization emerged like a cold dawn, **faster services don't matter if citizens don't trust them.**

By 2025, they hit the trust wall. The pilots worked. The scale-up stalled. And, leaders like Rony were discovering why.

The Questions That Kept Him Awake

Rony started asking simple questions about his Ferrari chatbot:

Where does citizen data reside? His team had built on a global cloud platform, which was fast, elegant, and powerful. But when he asked 'Where?' and 'Who?', the answers got uncomfortable. The data traveled across borders he couldn't control, through jurisdictions he couldn't govern.

Can decisions be explained? His chatbot served 2 million citizens. Feedback revealed a pattern, *it's fast, but... how did it decide? Where's my data going?*' He had no answers.

Who oversees this? The system ran on autopilot. No human reviews. No escalation paths. Just algorithms deciding people's lives.

Gartner's research echoed his fears:

- **61%** of citizens ranked secure data practices as extremely important
- But only **41%** trusted governments to protect their data
- And, **54%** wanted transparency into data use

Rony realized that speed without accountability isn't innovation. It's a risk.

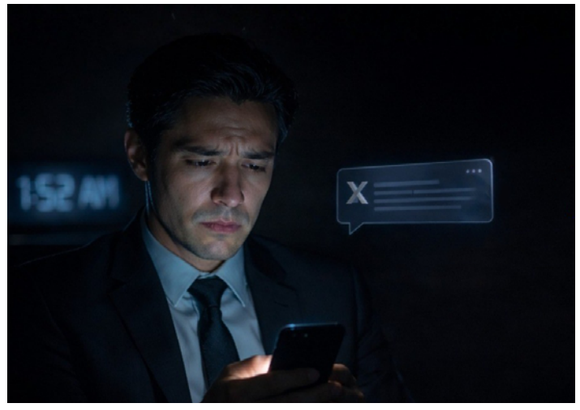

The Answers He Found

Rony discovered that forward-looking governments now follow **three pillars**:

1. Sovereign control of data: Citizen data stays within national boundaries, always. Not because the law demands it, but because citizens demand it.

2. Responsible AI operations: Decisions are explainable, bias-checked, and supervised by humans. The machine suggests; the public servant approves.

3. Unified digital platforms: No more bolting security onto systems built without it. Governance, automation, and service delivery work together from the start.

He learned that the right platform offers:

- **Governance-by-design**: Rules are embedded, not bolted on
- **Jurisdiction-aware data control**: Data knows where it belongs
- **Responsible AI decisioning**: Intelligence that can be explained and defended

Six Months Later

Rony's team deployed a new system. Same speed. Same scale. But now:

- Every decision carried an auditable trail
- Every data packet knew its jurisdictional boundaries
- Every citizen could request and receive an explanation

The chatbot still processes permits in 39 seconds. But now, Rony is not anxious and sleeps peacefully.

2026 marks digital government maturity. The race of 2024 is over. Governments that design services to be:

- **Fast by experience**: Citizens feel only the speed
- **Responsible by architecture**: Oversight built into every layer
- **Sovereign by default**: Data stays home, always

...will redefine how citizens interact with the state.

The question isn't whether your citizens will ask these questions. It's whether you'll have answers when they do.

Rony found his answers. Read the complete checklist, which has been designed to help you find yours, moving beyond the race of 2024 toward responsible, sovereign AI that you can actually defend.

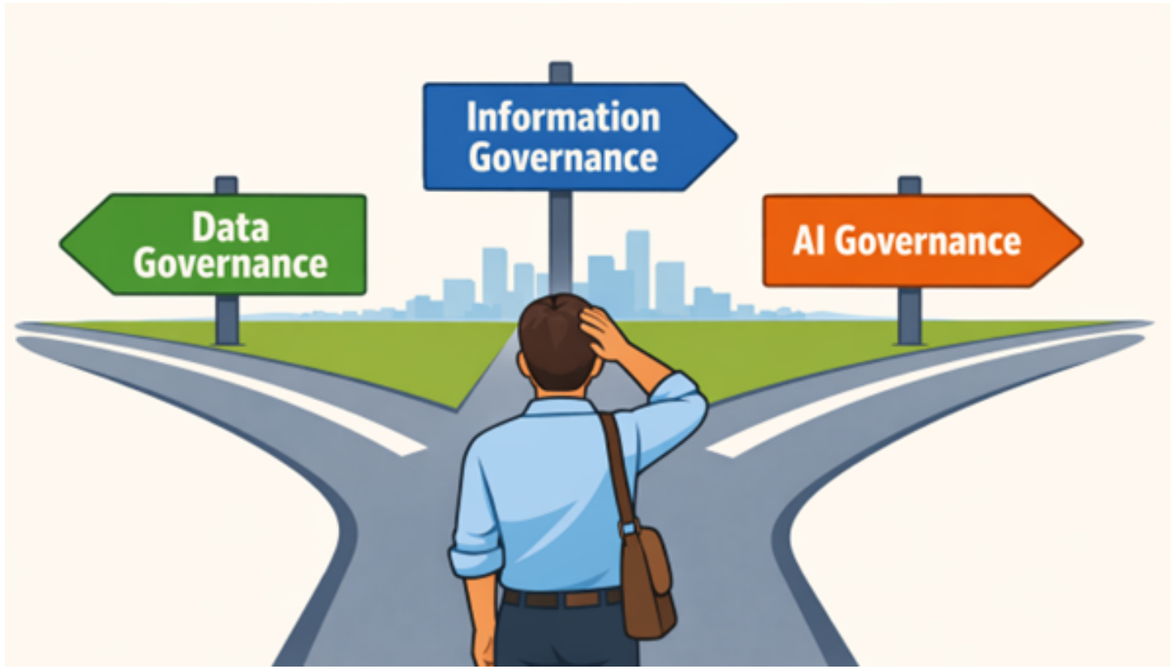

Why Data, Information, and AI Governance Can't Stay in Their Lanes

By Chad Barendse

If you've spent any time in governance, you've noticed something: we've been running three separate disciplines that are all trying to solve variations of the same problem.

Data governance over here. Information governance over there. And now AI governance is showing up like a new kid at school, expecting everyone to make room.

Each has its own framework. Its own teams. Its own language.

And somehow, we're surprised when the business doesn't know the difference between the three.

A Quick History Lesson

Governance didn't start with data. It started with records. Physical documents. Filing cabinets. Retention schedules. Making sure information was stored, retrievable, and destroyed when it needed to be.

Then digital happened. Data governance emerged to manage quality, integrity, lineage, and compliance across structured systems. Information governance continued in parallel — often with different reporting lines and different priorities.

For years, these worlds barely spoke to each other.

But the underlying intent was always the same: **make sure information is accurate, accessible, protected, and fit for purpose.**

We just kept creating new labels depending on where the information lived.

AI Doesn't Care About Your Org Chart

Here's what's changed.

AI models consume everything — structured data, unstructured documents, policies, emails, internal records. AI doesn't distinguish between a "data asset" and an "information asset." It just learns from whatever you feed it.

This is where things converge fast:

- **Data governance** provides the quality, lineage, and classification that AI needs to produce trustworthy outputs
- **Information governance** ensures that unstructured information — the stuff AI loves to consume — is properly managed, retained, and destroyed
- **AI governance** needs both of these foundations to answer the hard questions: What data trained this model? Is it classified correctly? Who's accountable? Should it have been destroyed already?

And here's what I want to challenge: **most AI governance principles aren't actually new.** We've been grappling with data ethics, privacy, and responsible use for years. AI governance extends those principles into machine learning and automated decision-making.

If you can't tell me what data trained your model and who's accountable for it — you don't have an AI governance problem. You have a **data and information governance problem that AI just made visible.**

Enterprise Knowledge Is Not Governance

One discipline that shouldn't be collapsed into governance is knowledge management.

Enterprise knowledge is becoming critical as AI adoption scales. Ontologies, semantic models, and shared business concepts are what turn raw information into something machines and humans can actually reason over.

But this is not governance work. It's data and knowledge management work.

Data and AI governance still have a role to play here. They set the expectations around standards, accountability, and risk for how enterprise knowledge is created and used. But the act of building and maintaining knowledge models? That belongs with the teams managing data and information — not the teams writing policy.

It's an important distinction. Governance sets the guardrails. It doesn't build the road.

The Convergence Is Already Happening

This isn't a theoretical prediction. It's playing out right now.

Look at the job market. I'm seeing more and more roles with **"Data & AI Governance"** in the title. Not separate positions — one person expected to span both disciplines. Organisations are already voting with their hiring decisions.

And it makes sense. The person governing your data quality, classification, and lineage is the same person who needs to understand what data is feeding your AI models, how it's being used, and whether it should even be there.

The skills overlap massively. The stakeholders are the same. The risks are connected.

The market is converging these roles faster than our frameworks are keeping up.

Do We Need One Framework or Two?

I'll be upfront, I don't think the answer is clear yet.

Maybe we end up with a single unified governance framework that covers data and AI together. Maybe we keep separate frameworks, a data governance framework and an AI governance framework but they integrate tightly and reference each other.

In practice, I think most organisations will land somewhere in the middle. You'll have your data governance policies covering quality, classification, lineage, and retention.

And you'll have AI governance policies covering model oversight, ethics, accountability, and responsible use. **But they'll need to work together as a system, not sit in separate folders on SharePoint.**

The AI governance framework can't exist without the data governance foundations underneath it. And the data governance framework needs to account for how data is being consumed by AI.

Separate documents? Maybe. Separate thinking? No.

Information Governance and Data Governance Should Merge

I don't see the distinction between information governance and data governance holding much longer. In many organisations, they're already being run together. The same teams handle classification, retention, destruction, metadata, and ownership across structured and unstructured information. The separation often exists in name and reporting lines rather than in day-to-day delivery.

The historical split made sense when data governance was only about structured data in databases and information governance handled everything else. But that line disappeared years ago. Modern data governance programs manage structured and unstructured data. They handle records retention and destruction. They deal with classification across all information types.

Some organisations will keep the disciplines formally separate for regulatory or organisational reasons. That's fine. But the direction of travel is clear. In practice, these capabilities are converging into a single way of governing information, regardless of format.

Where Does This Leave Us?

The way I see it, we're heading toward two core governance disciplines:

- **Data Governance** — covering all information assets, structured and unstructured, including what was traditionally information governance
- **AI Governance** — covering model oversight, ethics, accountability, and responsible use, deeply integrated with data governance

Two disciplines. Tightly connected. Possibly one team.

Enterprise knowledge isn't something governance owns — but it's something governance depends on.

The organisations that recognise this convergence early and build for it will move faster and govern more effectively. The ones still running three separate programs with three separate teams will keep wondering why things fall through the cracks.

Chad Barendse is Co-founder of DGX Group, a consultancy dedicated solely to Data and AI Governance. Originally published here

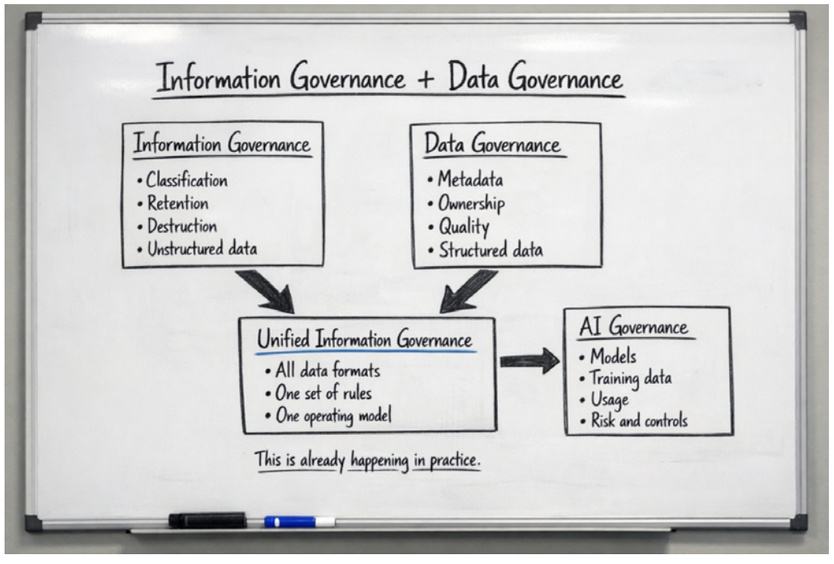

WA Local Governments Lag on Cyber

Western Australian local governments are failing to fix known IT security weaknesses, with 60 per cent of control flaws identified in 2025 carrying over unresolved from prior years.

A new report from WA Auditor General Caroline Spencer - Local Government 2025 - Information Systems Audit Results - found 333 control weaknesses across 68 local government entities.

Nine per cent of findings were rated significant, 69 per cent moderate and 22 per cent minor. Access management recorded the highest number of weaknesses, with 78 findings across 36 entities.

"These weaknesses put entities at greater risk of service disruptions, disclosure of ratepayers' data, financial loss and reputational damage," said Auditor General Spencer.

The report's case studies illustrate the consequences of poor controls. In one instance, a threat actor manipulated supplier account details in a financial system, resulting in a fraudulent payment of approximately $350,000 to an unknown third party.

Capability maturity assessments conducted at 15 entities showed an overall decline across all 10 control categories. Only one entity met the benchmark for access management and just two met the standard for endpoint security. Of the 11 entities assessed in both 2024 and 2025, results held steady in four categories but declined in six. The sharpest drops were recorded in risk management, change management, network security and access management.

Auditors found that one entity's internal corporate network was reachable from its public library, due to insufficient network segregation. Another entity left default administrator credentials unchanged on its building management system, exposing temperature, lighting and door controls to potential attack.

In the information security framework category, 54 per cent of entities lacked effective or up-to-date policies governing their IT environment. Half had no ICT steering committee to oversee technology strategy.

Human resource security weaknesses included insufficient phishing awareness training and failure to conduct police clearance checks on staff in privileged roles, such as finance officers and systems administrators.

Endpoint security failures were also prominent. Auditors found over 70 per cent of staff at one entity could run Microsoft Office macros - a known malware vector - due to misconfigured controls.

Physical security shortcomings included at least one server room lacking fire suppression systems and exhibiting structural damage to fire-rated walls, with old IT equipment stored inside.

The report also highlighted positive examples. One entity implemented continuous external security assessments covering firewall configuration, network penetration testing and Essential Eight mitigation strategies. Another deployed data loss prevention controls alerting to transfers of personally identifiable information.

The findings mirror concerns raised in the OAG's companion State Government 2025 - Information Systems Audit Results, tabled in December 2025, which identified persistent IT governance and cyber security weaknesses across WA state agencies.

Unstructured Data Outpacing Security Controls: CSA

Three-quarters of organisations are confident in their ability to secure unstructured data - yet more than two-thirds report less than 80% remains unprotected, a new industry survey has found.

The Rise in Unstructured Data and AI Security Risks report, published by the Cloud Security Alliance (CSA), reveals a widening gap between perceived and actual security posture across enterprise data environments. The survey polled 210 IT and security professionals in November 2025.

More than half (56%) of respondents reported only partial visibility into where their unstructured data is stored. A further 10% were unsure of their actual coverage.

Unstructured data - encompassing documents, emails, chat logs, images, and financial records - accounts for approximately 33% of enterprise data, with semi-structured data comprising an additional 21%. Nearly a third (29%) of organisations reported that unstructured data accounted for more than half of their annual data growth.

"The explosive growth of unstructured data - estimated by Gartner to account for between 70% and 90% of enterprise data - has become a defining characteristic of modern organisations," said Hillary Baron, AVP of Research, Cloud Security Alliance.

"What is clear from this study is that as unstructured data continues to expand and spread across environments, many organisations are struggling to keep pace with the visibility, governance, and protections needed to manage it securely."

Despite listing security (74%), governance (57%), privacy (54%), and compliance (50%) as their top concerns, organisations are failing to execute foundational controls.

Nearly one-third (32%) of organisations use 11 or more tools to manage unstructured data. Of those, 12% rely on at least 21 tools. Data encryption (62%), cloud security (60%), application security (59%), and identity and access management (56%) are among the most widely deployed.

The full report is available for download here.

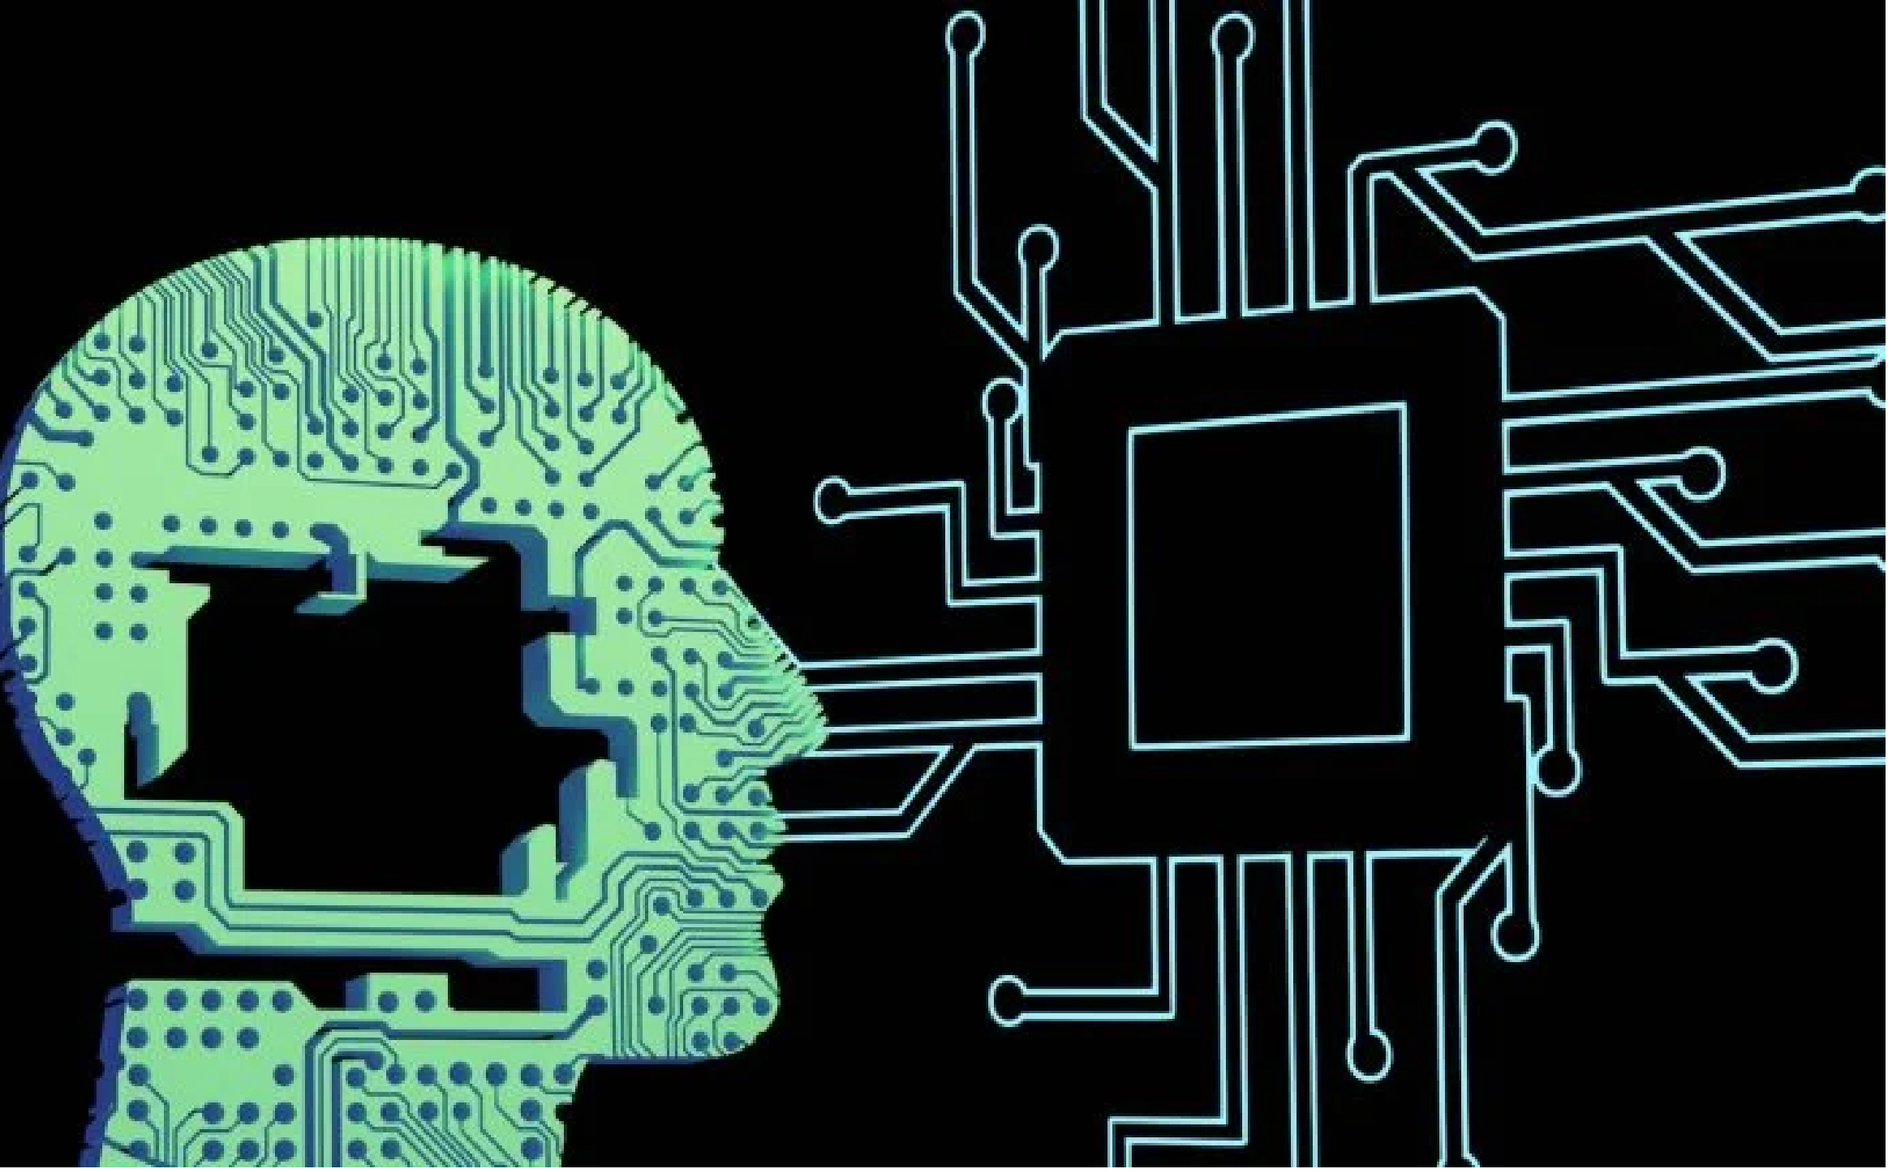

AI Automation Arrives as Rising Tide, Not Crashing Wave, MIT Study Finds

Artificial intelligence is advancing across the full range of text-based knowledge work simultaneously, rather than disrupting narrow task categories in sudden surges, according to preliminary research published by MIT FutureTech in April 2026.

The study, *Crashing Waves vs. Rising Tides: Preliminary Findings on AI Automation from Thousands of Worker Evaluations of Labor Market Tasks*, analysed more than 17,000 expert evaluations of AI responses across 3,000-plus tasks drawn from the US Department of Labor O*NET database.

Researchers found AI success rates - measured as responses a manager would accept without editing - averaged 60 per cent across all models tested. Success rates are projected to reach 80 to 95 per cent on most text-based tasks by 2029.

The report states that AI capabilities are "already substantial and poised to expand broadly," with progress resembling "a rising tide, with widespread gains across many tasks simultaneously," rather than concentrated disruption of specific task categories.

The research estimates that the maximum task duration AI can handle at a 50 per cent success rate has been doubling every 3.8 months since mid-2024. By Q3 2025, frontier models were achieving 50 per cent success rates on tasks taking humans up to one week to complete.

Failure rates are also dropping rapidly. The study calculates that failure rates halve every 2.4 to 3.2 years across tasks taking five minutes to 24 hours of human effort. Over the 2024-Q2 to 2025-Q3 observation window, this translated to annual success rate improvements of 8 to 11 percentage points.

The study also found that newer model releases - rather than simply larger models - drive the most consistent performance gains. Newer models improve performance roughly equally across tasks of all durations, while larger models show greater gains on short tasks but diminishing returns on longer ones.

Among the job families studied, Computer and Mathematical roles had the highest proportion of LLM-automatable tasks, at 94 per cent. Business and Financial Operations followed at 93 per cent, and Management at 91 per cent.

Legal tasks (85 per cent automatable by LLMs), Life, Physical and Social Science (85 per cent), and Architecture and Engineering (84 per cent) also ranked highly. Healthcare Practitioners and Technical roles showed 62 per cent of tasks with meaningful LLM time-saving potential.

Average AI success rates varied substantially by domain. Legal tasks achieved only 47 per cent success - the lowest

Large-scale worker evaluation of 3,000-plus tasks finds LLM capabilities advancing broadly across knowledge work - giving organisations time to plan, but no time to delay.

of any category studied. Installation, Maintenance and Repair recorded the highest at 73 per cent, though the study restricted its focus to the text-based components of each role.

The researchers caution that high task-level success rates do not directly translate to equivalent job automation. "Last-mile" implementation costs, data integration challenges, regulatory constraints, and organisational adoption timelines mean real-world deployment will lag AI capability growth.

What Rising Tides Means for Planning

The study argues that the gradual, broad-based nature of AI improvement is strategically significant. The authors write that "workers are likely to have some visibility into these changes, rather than facing discontinuous jumps in AI-driven automation."

However, the authors also warn that achieving near-perfect AI performance - success rates approaching 100 per cent - will take considerably longer than reaching the 80 to 95 per cent threshold. Tasks with low tolerance for errors, such as legal compliance and records management, will require sustained human oversight well beyond 2029.

The report notes that the pace of improvement documented "implies that LLMs will be able to complete most text-related tasks with success rates of, on average, 80-95% by 2029 at a minimally sufficient quality level." Achieving near-perfect success rates or comparable rates at superior quality "would require several additional years."

The research explicitly distinguishes between task automation and workforce impact, citing work by economists David Autor and Neil Thompson that automation of individual tasks does not necessarily reduce wages or employment in a given occupation. Labour market impacts depend on how entire occupation task bundles respond to automation.

Methodology and Limitations

The study evaluated 41 AI models released between June 2023 and August 2025, including models from Anthropic, Google, Meta, Mistral, OpenAI, and others. Responses were rated by domain experts recruited via the Prolific platform, all with at least six months of relevant on-the-job experience.

Approximately 34.6 per cent of collected responses were excluded after quality assurance checks, including attention checks and response-consistency screening. The authors note the survey is ongoing and current results are preliminary.

The paper contrasts its findings with prior benchmark-based research by METR, which reported steeper capability curves consistent with a "crashing wave" pattern. The MIT team attributes the divergence partly to the use of real-world, non-deterministic labour market tasks rather than deterministic software-engineering benchmarks.

The full paper is available at https://arxiv.org/abs/2604.01363.

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

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

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

Knowledge Unlocked with AFIVE AI

Fragmented enterprise knowledge stores have a new challenger. Sydney-based IT services firm Adactin has released AFIVE, an AI-powered platform that uses Retrieval-Augmented Generation (RAG) to give employees natural language access to documents held across multiple cloud repositories.

Built on Microsoft Azure OpenAI Service and Azure AI Foundry, AFIVE connects to SharePoint, Google Drive, Azure Blob Storage and Dropbox.

Users query the platform in plain language and receive synthesised responses drawn from those sources, without needing to search each system separately.

The platform uses LangChain to co-ordinate the RAG pipeline, with vector databases providing semantic document retrieval.

Integration with Power BI and Power Automate allows the platform to surface insights and trigger automated workflows from within the same interface.

Security controls include role-based access via Microsoft Entra ID, network isolation and encrypted credential management.

Adactin says new knowledge repositories can be added without disrupting existing workflows, and the architecture is designed to scale as organisational data volumes grow.

"By automating repetitive retrieval tasks and enabling real-time summarisation, AFIVE redirects employee effort from administrative searching to higher-value strategic work", says Srinivas Gutta, Technical Practice Director, Adactin.

"In addition, by centralising organisational knowledge and enabling intelligent search, the platform reduces cognitive load, enhances collaboration, and increases confidence in decision-making across the enterprise.

The announcement reflects a broader push by mid-market IT services firms to productise AI capabilities developed for client engagements.

RAG architecture has become a common pattern for enterprise AI deployments because it allows large language models to draw on proprietary data without requiring retraining, reducing both cost and data exposure risk.

Founded in Sydney in 2011, Adactin employs more than 300 professionals across the APAC region and focuses on digital transformation, Microsoft and AWS ecosystem services, and AI-enabled engineering.

https://www.adactin.com

No-Code AI Agent to Connect Document Workflows

The friction of connecting AI document processing platforms to enterprise business systems has long stalled automation projects in regulated industries. A new tool from Melbourne-based Affinda aims to remove that barrier, allowing organisations to describe integrations in plain language and have the software write the code automatically.

Affinda's AI Integration Agent connects the company's Intelligent Document Processing (IDP) platform with downstream business systems - the vendor claims compatibility with more than 2,800 applications - without requiring developers or custom API builds.

The agent writes, refines, and tests integration code based on plain-language instructions provided by users. Among the supported enterprise platforms are Xero, Excel, Salesforce, Hubspot, Dynamics 365, OneDrive and Power Automate.

The tool targets two distinct audiences: organisations with in-house developers who want to rapidly prototype automation workflows, and smaller organisations without technical staff that have previously found integration costs prohibitive.

Affinda Head of AI Andrew Bird said the agent allows organisations to automate data export to "virtually any system - whether it's ERP, CRM, or other databases - using just natural language instructions."

Integration complexity has been a persistent barrier in IDP deployments. Organisations have typically faced a choice between rigid, pre-configured connectors or expensive custom API development - both of which add time and cost to automation projects before any documents are actually processed.

The Integration Agent attempts to address this by generating the integration code itself.

Affinda General Manager Charlie Bellingham cited a specialist lending firm as the type of organisation previously locked out of automated document and records workflows.

"The set-up time is drastically reduced from days or weeks - to something that can be done in 15 minutes," Bellingham said.

The broader context is a competitive shift in the IDP market. Traditional machine learning approaches required hundreds of training documents and ongoing retraining cycles.

Newer platforms - including Affinda - use large language models (LLMs) and retrieval-augmented generation (RAG) to reduce configuration time and handle greater document variability.

While this lowers the technical barrier to adoption, it also introduces new considerations around prompt governance and the traceability of AI-driven decisions in regulated environments.

"These are considerations core to our offering. Every output is traceable back to the source document, so data are auditable and defensible – which is exactly what regulated environments demand," Bellingham adds.

https://www.affinda.com

AI Tool Tames Spreadsheet Data Intake

	C	D	E	F	G	H	I	J	K	L	M	
it	Currency	Year	Scenario	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	
	USD	2022	Actuals	$78,338,286	$58,034,151	$75,874,351	$87,725,417	$84,878,696	$73,990,027	$95,372,852	$83,104,256	$5
	USD	2022	Actuals	($32,377,742)	($26,379,976)	($33,823,729)	($41,186,038)	($40,605,218)	($31,555,193)	($46,768,008)	($33,951,735)	($2
	USD	2022	Actuals	($3,477,740)	($2,612,821)	($3,043,000)	($4,206,639)	($3,635,596)	($3,390,988)	($4,751,711)	($3,385,162)	($
	USD	2022	Actuals	($8,777,061)	($5,956,476)	($8,788,853)	($9,995,832)	($8,550,645)	($8,956,466)	($11,183,485)	($8,871,134)	($
	USD	2022	Actuals	($851,217)	($648,182)	($827,626)	($971,146)	($1,009,775)	($865,150)	($1,184,627)	($992,651)	(
	USD	2022	Actuals	($3,158,339)	($2,728,823)	($3,421,161)	($4,332,736)	($4,162,923)	($3,139,829)	($4,315,056)	($3,425,707)	($
	USD	2022	Actuals	($4,069,215)	($3,519,759)	($4,009,533)	($4,836,632)	($5,257,809)	($4,412,376)	($4,985,773)	($4,979,583)	($
	USD	2022	Actuals	($6,585,349)	($4,759,687)	($6,493,370)	($6,546,763)	($6,734,243)	($6,212,732)	($8,195,799)	($6,417,083)	($
	USD	2022	Actuals	($1,828,234)	($1,342,350)	($1,751,692)	($2,175,232)	($1,829,981)	($1,789,035)	($1,954,312)	($2,007,974)	($
	USD	2022	Actuals	$19,584,572	$14,508,538	$17,451,101	$22,808,608	$21,219,674	$19,977,307	$21,935,756	$24,931,277	$1
	USD	2022	Actuals	($8,362,230)	($6,091,621)	($8,316,059)	($10,159,328)	($8,624,806)	($9,693,159)	($9,579,056)	($9,973,478)	($
	USD	2022	Actuals	($915,086)	($612,941)	($823,113)	($916,944)	($880,691)	($867,181)	($960,417)	($1,047,426)	(
	USD	2022	Actuals	($2,065,945)	($1,474,640)	($2,117,096)	($2,790,604)	($2,540,908)	($2,325,558)	($2,387,390)	($3,056,599)	($
	USD	2022	Actuals	($240,447)	($157,072)	($183,016)	($241,630)	($239,007)	($238,586)	($266,862)	($274,476)	(
	USD	2022	Actuals	($796,397)	($610,923)	($729,783)	($913,075)	($1,012,763)	($882,471)	($1,093,311)	($1,137,735)	(
	USD	2022	Actuals	($1,199,233)	($811,310)	($900,254)	($1,164,976)	($1,151,854)	($1,034,607)	($1,285,049)	($1,518,350)	(
	USD	2022	Actuals	($1,445,011)	($1,126,211)	($1,306,570)	($1,839,960)	($1,539,232)	($1,482,785)	($1,588,028)	($1,880,876)	($
	USD	2022	Actuals	($420,494)	($304,699)	($395,657)	($479,567)	($522,262)	($404,189)	($497,618)	($499,475)	(
	USD	2022	Actuals	$30,551,932	$19,731,611	$24,279,792	$35,090,167	$26,312,396	$25,896,509	$31,473,041	$27,424,404	$2
	USD	2022	Actuals	($13,819,091)	($9,286,529)	($10,933,750)	($15,827,962)	($12,543,885)	($11,801,803)	($15,410,493)	($12,281,930)	($
	USD	2022	Actuals	($1,369,961)	($942,250)	($1,067,439)	($1,630,125)	($1,266,063)	($1,202,899)	($1,418,412)	($1,366,956)	(
	USD	2022	Actuals	($3,344,503)	($2,300,612)	($2,595,976)	($3,778,056)	($3,167,287)	($2,947,517)	($3,546,969)	($2,816,091)	($
	USD	2022	Actuals	($338,744)	($244,628)	($271,723)	($354,541)	($300,962)	($260,911)	($327,018)	($336,207)	(
	USD	2022	Actuals	($1,468,744)	($973,887)	($1,032,242)	($1,717,271)	($1,287,760)	($1,206,283)	($1,470,777)	($1,246,641)	($
	USD	2022	Actuals	($1,633,292)	($1,089,178)	($1,472,323)	($2,057,777)	($1,473,412)	($1,538,962)	($1,683,014)	($1,385,482)	($
	USD	2022	Actuals	($2,276,556)	($1,464,050)	($2,030,413)	($2,759,580)	($2,279,559)	($2,087,858)	($2,509,070)	($2,137,231)	($
	USD	2022	Actuals	($661,654)	($401,978)	($494,551)	($729,508)	($547,250)	($536,085)	($699,474)	($582,080)	(
	USD	2023	Budget	$89,862,727	$99,687,807	$99,378,570	$55,271,910	$83,758,431	$51,637,163	$76,348,472	$66,086,281	$8
	USD	2023	Budget	($39,040,130)	($39,921,367)	($45,671,498)	($25,801,859)	($39,706,126)	($24,992,558)	($31,073,960)	($27,470,867)	($3
	USD	2023	Budget	($4,014,208)	($4,730,507)	($4,051,116)	($2,388,743)	($3,788,314)	($2,165,366)	($3,369,016)	($2,871,769)	($
	USD	2023	Budget	($11,164,524)	($11,563,467)	($10,590,962)	($6,628,891)	($10,201,147)	($6,336,226)	($8,308,362)	($8,049,766)	($

Fragmented spreadsheet data flowing into enterprise systems without standardisation is a chronic pain point for compliance-heavy industries - one that AI vendor Fisent Technologies says it has tackled with a new capability called Tabulate, added to its BizAI platform.

The tool claims to convert inconsistently formatted spreadsheets into structured, clean data without requiring middleware or extensive manual post-processing.

Tabulate is the sixth Agentic Action added to its Fisent BizAI agentic solution, joining Classify, Split, Extract, Verify, and Analyse. According to Fisent, it embeds data transformation logic directly into the extraction workflow rather than relying on large language model (LLM) inference for each operation - an approach the company says reduces processing costs and improves consistency for routine tasks.

Financial documents - invoices, policy schedules, loan records - routinely arrive in varied formats across different business systems. A well-documented constraint of current LLMs is their context window limit, which makes processing multi-thousand-row spreadsheets unreliable. Outputs often require significant manual correction before they can feed into automated downstream processes.

Fisent says Tabulate addresses this by handling common format normalisation tasks - such as converting date variants like “January 15, 2026” and “01/15/26” into a single standard - through direct transformation operations rather than LLM inference.

The capability also identifies correct table boundaries in spreadsheets with multiple tables and inconsistent layouts and removes formatting artefacts before data is passed to downstream systems.

Fisent says Tabulate is already deployed in production environments for customers in insurance underwriting and commercial lending. Financial services customers are also reportedly using the tool to prepare data for financial crime and risk management reporting in cloud data warehouses, including Snowflake.

The company describes its broader BizAI platform as “production-proven” with Fortune 500 customers in banking, insurance, and wealth management.

The challenge of preparing unstructured and inconsistently formatted data for AI-driven automation is well recognised across enterprise deployments. As organisations extend into agentic AI - where software takes autonomous action based on extracted data - data quality, consistency, and auditability become critical compliance and governance concerns.

Regulators in banking, insurance, and financial services sectors increasingly expect demonstrable controls over data provenance in automated decision-making systems.

https://www.fisent.com

Partnership Targets Insurance Compliance

Insurance technology provider ICE-Tech has integrated intelligent document processing technology from TCG Process into its Alice software platform, automating the ingestion, validation and routing of unstructured content — including emails and attachments — across policy issuance, claims and mid-term adjustment workflows.

ICE-Tech’s Alice framework structures AI deployments around three published requirements: Compliant, Open and Safe. Under this model, any partner solution integrated into Alice must demonstrate explainability, a full audit trail and human-in-the-loop controls before deployment in the ICE-Tech cloud environment.

The framework’s “Open” designation refers to architecture openness — partner solutions are permitted provided they meet ICE-Tech’s compliance and safety thresholds — rather than open-source licensing.

TCG Process’s document ingestion technology was assessed against what the announcement describes as “AI/ML safety standards.”

TCG Process’s core product, DocProStar, is an intelligent document processing (IDP) solution that extracts and classifies data from unstructured sources. Within the Alice framework, this capability is applied to the specific document flows common in insurance — FNOL emails, policy documents, attachments and correspondence — and routes extracted data into downstream policy and claims systems.

The automation targets the re-keying burden that arises when claims or policy teams manually extract information from incoming documents before entry into core systems.

The integration is already live in production with an unnamed client described by both vendors as a “leading specialist insurer.”

The reported outcomes include:

- More than 30,000 manual interactions removed per year
- 30 to 40 minutes saved per claim through automated First Notice of Loss (FNOL) intake
- A claimed 75% increase in claims operational input capacity
- Exception handling time reduced from 15-25 minutes to an average of four minutes

https://ice-tech.com https://www.tcgprocess.com

Document Automation without the Uplift

Automated redaction, natural-language AI agent creation, and realtime compliance reporting are among a new wave of updates Hyland has announced for its Content Innovation Cloud platform and associated enterprise content management (ECM) products.

Hyland’s emphasis on federation - enabling AI to operate across repositories without migration - responds to the practical reality that most large organisations manage content across multiple systems accumulated over decades.

The new updates span six products - Hyland Automate, Hyland Intelligent Document Processing (IDP), Content Federation Service, OnBase, Alfresco, and Nuxeo - and are aimed at organisations seeking to automate document-intensive compliance workflows without migrating data between systems.

The automated redaction capability within Hyland Automate targets the removal of sensitive information from documents as part of workflow processing - a capability relevant to privacy obligations under frameworks including the Australian Privacy Act and sector-specific regulations in government, healthcare, and financial services.

Hyland Automate also now allows users to create and deploy AI agents using natural-language prompts, a capability Hyland says reduces the technical barrier to workflow automation.

Hyland IDP, the company’s intelligent document processing product, gains AI-driven document classification and expanded file and text recognition support in this release.

A new reporting module provides realtime visibility into processing accuracy, throughput, and compliance metrics - an addition likely to appeal to governance and risk managers who require audit trails for automated document handling.

The Content Federation Service update extends connectors to include SharePoint 365 alongside existing OnBase, Alfresco, and Nuxeo repositories, allowing content to be accessed and processed across systems without physical migration.

For organisations managing records across multiple legacy platforms - a common situation in government agencies and large enterprises - this approach avoids the cost and risk of full data migration while enabling AI-assisted search and processing across disparate content stores.

Platform-level updates to Alfresco include enhanced security controls and revised authorisation flows, along with expanded developer tooling.

https://www.hyland.com/en

Agent Supervision via Ping Identity

Enterprises deploying autonomous AI agents have a governance problem that traditional identity and access management was never designed to solve - and identity vendor Ping Identity is targeting that gap with a new product suite aimed at controlling what AI agents do at the moment they act.

Ping Identity has announced General Availability of Identity for AI, comprising three components: Agent IAM Core, Agent Gateway, and Agent Detection.

Agent IAM Core onboards, authenticates, and authorises AI agents as a distinct identity type with defined ownership and policy.

Agent Gateway provides a runtime enforcement layer for agent-to-system interactions, centralises monitoring and audit, and supports the Model Context Protocol (MCP) - an emerging integration standard that allows AI agents to connect directly to applications and data sources without a human intermediary.

Agent Detection, delivered via PingOne Protect, uses behavioural signals and bot authentication protocols to identify AI agents in operation and feed risk signals into authorisation decisions.

The product addresses a structural limitation in traditional IAM architecture. Conventional access management governs who can log in and what they can access at that point, but does not enforce controls over what authenticated agents do thereafter.

Identity for AI replaces agent impersonation of human credentials with delegated, scoped tokens, applying least-privilege enforcement at each individual action rather than at the point of authentication. Human accountability is retained under the model, with approval workflows supported for both on-behalf-of-user and fully autonomous flows.

MCP support is a significant element for enterprises evaluating AI agent infrastructure. The protocol allows AI agents to retrieve data, trigger workflows, and act inside enterprise systems at machine speed without routing work through a human. Security researchers have flagged poorly governed MCP connections as high-value attack paths, noting that compromised MCP tokens could allow threat actors to operate within critical systems with minimal visibility. Ping's Agent Gateway is positioned to standardise enforcement across MCP-based integrations without requiring organisations to rewrite existing services.

"AI agents are not features. They are actors in the enterprise that require identity, authority, and accountability," said Andre Durand, CEO and Founder of Ping Identity.

"Identity is foundational. Agents acting autonomously at agentic scale and speed against systems of record will require continuous verification and enforcement at every decision."

Ping is not alone in targeting this space. Microsoft announced its Agent 365 platform - described as a control plane for AI agents incorporating Microsoft Defender, Entra, and Purview capabilities - for General Availability on 1 May 2026.

Entro Security launched Agentic Governance and Administration (AGA) at RSA Conference 2026, and Veza has introduced Access Agents targeting AI identity governance.

A Cloud Security Alliance and Oasis Security report found that 78% of organisations lack formal policies for creating or removing AI identities, and 92% are not confident their existing IAM platforms can handle AI agents.

https://www.pingidentity.com

Governance Platform Detects AI Agents

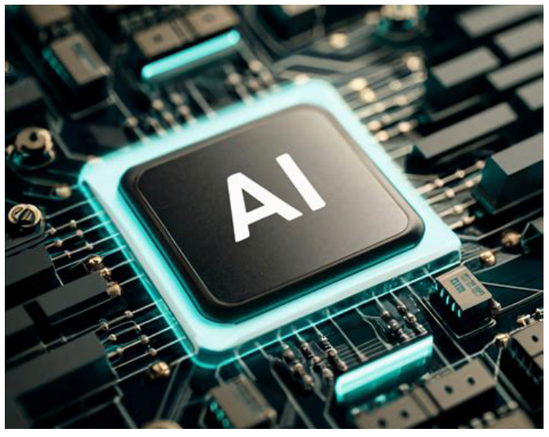

Realtime monitoring and automated guardrail enforcement are the foundations of a significant expansion to OneTrust's AI governance platform, aimed at organisations struggling to keep compliance controls pace with rapidly scaling AI deployments.

The US-based vendor has added three capabilities to its platform: AI Agent Detection and Inventory, a Policy Manager and Policy Library, and AI Guardrail Enforcement. The additions are designed to shift AI governance from periodic, point-in-time compliance reviews to continuous oversight across agents, models, and datasets.

The move reflects growing pressure on governance, risk, and compliance teams to account for AI systems that were not covered by traditional IT asset management. The introduction of agentic AI - autonomous systems that can take actions, access data, and interact with other systems with minimal human oversight - has exposed gaps in most organisations' existing risk frameworks.

The EU AI Act, which began applying obligations to high-risk AI systems from August 2024, and the US National Institute of Standards and Technology (NIST) AI Risk Management Framework have both heightened demand for structured AI governance tooling.

Australia's government has released voluntary AI Ethics Principles but has signalled interest in mandatory guardrails for high-risk applications, adding urgency for enterprises operating across jurisdictions.

What the New Capabilities Do

The Agent Detection and Inventory capability is designed to continuously discover and catalogue AI agents, models, and datasets across an organisation's environment, capturing ownership, data access, and lineage information.

The intent is to eliminate blind spots in AI asset registers - a persistent problem as development teams adopt no-code and low-code AI tooling that can bypass traditional procurement or IT review processes.

The Policy Manager includes a library of pre-built policies aligned to frameworks including the NIST AI RMF and the EU AI Act. The platform translates those framework requirements into monitored controls, with evidence capture intended to support audit readiness. OneTrust claims that the system provides realtime visibility as AI systems evolve.

Guardrail Enforcement targets generative AI, traditional machine learning models, and agents. The system is described as inspecting AI systems continuously and automatically enforcing protections - such as blocking or limiting personal data exposure

The platform integrates with Amazon Bedrock, Amazon SageMaker, Azure AI Foundry, Azure OpenAI, Databricks Unity Catalog, and Google Vertex AI. This breadth of integration is significant for enterprise teams that run AI workloads across multiple cloud environments, a common scenario in large government agencies, financial institutions, and healthcare organisations.

OneTrust operates in a growing AI governance market that includes competitors such as ServiceNow, IBM OpenPages, and emerging specialist vendors. Gartner has flagged AI governance tooling as a priority area for 2025-26 as regulatory obligations across jurisdictions increase.

https://www.onetrust.com

OpenText Unified Security Suite

With ransomware hitting 42% of Australian businesses in the past year and new cybersecurity legislation now in force, OpenText has localised its Secure Cloud platform for Australian managed security providers (MSPs), adding AUD billing, automated provisioning, and analytics designed to help partners identify and close client security gaps.

The platform consolidates endpoint protection, email threat and DNS protection, security awareness training, email encryption, managed detection and response (MDR), endpoint detection and response (EDR), and cloud backup into a single multi-tenant environment. It is targeted at MSPs serving small-to-medium business and mid-market clients.

The platform promises a single pane of glass for security and data protection, with reporting and analytics built in.

The Australian launch coincides with two significant pieces of legislation that are reshaping data and cybersecurity obligations for organisations and their service providers. The Cybersecurity Act 2024 introduced mandatory security standards and incident reporting obligations, while the Privacy and Other Legislation Amendment Act 2024 strengthened data protection requirements, including stricter rules around how personal data is stored, managed, and protected across networks, devices, and cloud platforms.

Steve Stavridis, Regional Vice President APAC at OpenText Cybersecurity, said the company was seeing a clear shift among Australian partners "toward managed security services and hybrid delivery models."

He added that organisations' data privacy practices were "being tested like never before," with customers demanding stricter controls over data storage and management.

MSPs operating under these frameworks will need to demonstrate their platforms support client compliance. Organisations relying on MSPs for security and backup services should consider how platform-level reporting from tools like Secure Cloud maps to their own governance and audit requirements.

According to OpenText Cybersecurity's own 2025 Global Managed Security Survey, 42% of Australian businesses reported a ransomware compromise in the past year. More than half experienced multiple attacks. The survey found 82% of affected businesses were unable to recover all encrypted or stolen data, and 34% admitted paying a ransom.

The Australian Signals Directorate's Annual Cyber Threat Report has consistently ranked ransomware among the most significant threats to Australian organisations, noting the increasing targeting of critical infrastructure and government entities.

https://cybersecurity.opentext.com/

Security vendor bets on semantic AI to tame agent risk

A governance gap is stalling enterprise AI deployments - and data security vendor Rubrik says it has a fix built on a custom language model.

The company has launched SAGE (Semantic AI Governance Engine) at RSAC 2026 in San Francisco. Rubrik describes it as the data security industry's first AI governance engine for controlling autonomous agents in real time. SAGE forms the core of Rubrik Agent Cloud, the company's platform for deploying and auditing AI agents at enterprise scale. It is driven by a proprietary Small Language Model (SLM) that Rubrik developed in-house.

The underlying problem SAGE targets is well recognised. Existing enterprise governance frameworks rely on deterministic, keyword-based rules. Those rules cannot interpret natural-language instructions or respond dynamically when an AI agent takes an unexpected action.

Rather than matching keywords, SAGE interprets the intent behind a written policy. An instruction such as "Do not give financial advice" is converted into machine-enforceable logic that the engine can apply to ambiguous or novel agent behaviours.

Announced capabilities include adaptive policy improvement, which Rubrik says flags ambiguous guardrails before a violation occurs, and a remediation feature called Agent Rewind. Rubrik claims Agent Rewind can instantly undo destructive agent actions and restore data integrity.

Rubrik conducted an internal benchmark comparing its SLM against OpenAI's GPT-5.2 for policy violation detection. The company reported its model processed messages five times faster and detected violations more accurately.

https://www.rubrik.com

DocSolid Automates Metadata for Mass Records Digitisation

A new capability from law firm records software vendor DocSolid allows scanning operators to automatically retrieve metadata from existing records management systems during large-scale physical records digitisation - without requiring direct access to those systems.

Records Connect, an extension of DocSolid's Airmail2 Cloud high-volume scanning platform, uses barcode identifiers already printed on physical files and folders to query connected records management systems and auto-populate document profiles during scanning. The company says it integrates with iManage IRM, FileTrail and Microsoft SharePoint.

The announcement targets a longstanding friction point in backfile digitisation: manual metadata entry. In large-scale projects involving millions of physical records, operators have historically needed familiarity with document management systems (DMS) to correctly profile each item. Records Connect is designed to bypass that requirement by pulling existing metadata automatically.

For records managers and information governance professionals overseeing digitisation programmes, the capability speaks to a broader challenge: maintaining metadata integrity and retention schedule compliance when transitioning from physical to digital records at scale.

DocSolid says Records Connect supports both live database connections and data warehouse integrations, ensuring scanning teams access current records data.

The product also addresses outsourcing constraints. Firms increasingly rely on third-party scanning operators for backfile conversion, but providing external staff with DMS credentials or training carries security and compliance risks. Records Connect is positioned as a way to extend profiling capability to outsourced labour without granting system access.

David Guilbault, DocSolid's Vice President of Customer Experience, said the capability "dramatically" improves efficiency on large digitisation projects.

https://www.docsolid.com

Governance Tool for Enterprise Automation

Growing governance gaps in enterprise automation environments have prompted US-based observability vendor Reveille Software to launch a dedicated enterprise platform tier, targeting organisations managing complex document processing and AI-driven workflow deployments.

The move reflects deepening industry concern that monitoring tools designed for simpler IT environments are insufficient when robotic process automation (RPA), enterprise content management (ECM), and intelligent document processing (IDP) platforms underpin regulated business processes.

The new product, Reveille Enterprise, restructures what was previously a single-tier offering into two distinct editions. Reveille Standard continues to serve smaller or less complex environments focused on operational visibility, while the enterprise edition adds governance, accountability modelling, and self-healing automation for large, distributed, or regulated organisations.

For records managers, compliance officers, and information governance professionals, the most significant addition is a role-based ownership and accountability model linked to SLA-driven event management.

This addresses a recognised challenge in automation deployments: when automated processes span multiple platforms and teams, accountability for service failures is often unclear.

The vendor says the platform enables organisations to define ownership at a granular level and trigger escalations when SLA thresholds are breached.

Additional capabilities in the enterprise tier include integration with IT service management (ITSM) and AIOps platforms, support for the OpenTelemetry observability standard, user experience analytics for identifying real-user anomalies, and - according to Reveille - more than 90 automated remediation actions to reduce mean time to recovery (MTTR).

The company also promotes a distributed server architecture with remote collectors and centralised administration for large environments. These figures are vendor claims and have not been independently verified.

The platform adds support for the Model Context Protocol (MCP), an open standard introduced by Anthropic in November 2024 that is increasingly adopted as a universal interface for connecting AI agents to enterprise tools and data sources.

OpenAI, Google DeepMind, and Microsoft have since adopted the standard, and it was transferred to the Linux Foundation-backed Agentic AI Foundation in December 2025.

https://www.reveille.com

On-Premise AI Stack Targets Regulated Industries

AI knowledge management vendor Docsie has released a fully on-premise platform that runs entirely on customer-owned hardware and customer-controlled language models. The company says it is responding to a growing unwillingness among regulated enterprises to route sensitive data through third-party cloud infrastructure.

The platform - branded as a Bring-Your-Own-Model (BYOM) stack - connects existing large language model (LLM) endpoints, including Llama, Qwen, DeepSeek, Mistral, and any OpenAI-compatible model, to document management and compliance workflows. Docsie says no data leaves the customer's network at any point.

The timing reflects a widening gap between AI investment and AI deployment. According to Info-Tech Research Group's AI Trends 2026 report. 72% of global IT leaders now list data sovereignty and regulatory compliance as their top AI-related challenge - up from 49% the previous year.

A separate Nutanix 2026 Enterprise Cloud Index survey found 80% of respondents consider data sovereignty a high priority when making infrastructure decisions.

The platform targets organisations in financial services, healthcare, manufacturing, defence contracting, and ERP implementation - sectors where regulatory frameworks routinely prohibit sensitive data from transiting external networks.

"The enterprise AI problem isn't capability - it's control," said Philippe Trounev, CEO of Docsie.

"These organisations have bought GPUs, they're running models internally, they have vLLM or Bedrock deployed - but they have no way to connect that inference capacity to their actual knowledge management workflows."

A content compliance scanning module analyses video, audio, and text against personally identifiable information (PII), brand, and custom policy frameworks. Docsie says an interactive timeline viewer allows compliance officers to click flagged violations and jump to the corresponding moment in training video - including content visible on screen for less than a second.

An air-gapped documentation delivery function packages complete offline documentation sets deployable to disconnected networks, factory floor terminals, or field equipment with no internet dependency. The package ships as a Docker container for any Kubernetes cluster.

A multi-agent orchestration layer allows organisations to build domain-specific AI agents - described in release materials as compliance reviewers, standard operating procedure (SOP) generators, and training content converters - that connect to enterprise systems including Jira, Salesforce, and ServiceNow. Docsie says no code is required to configure these agents.

The platform includes a multi-tenant enterprise portal with single sign-on (SSO), per-organisation vector indexes, and session-level audit trails with remote revocation - features relevant to organisations managing access governance across multiple business units or client environments.

Docsie is an AI-powered knowledge orchestration platform that converts training videos, PDFs, and existing content into structured knowledge bases, then delivers them as branded portals with AI chat, compliance scanning, and learning management.

https://www.docsie.io

Ungoverned AI a Growing Compliance Risk

When AI agents start taking actions rather than merely answering questions, the integrity of their knowledge source becomes a compliance problem. Knowledge management vendor eGain has released a set of platform connectors designed to anchor Microsoft Copilot, Anthropic Claude, Google Gemini CLI, and the Cursor developer environment to a single governed knowledge repository.

The connectors link those AI platforms to eGain's AI Knowledge Hub via the Model Context Protocol (MCP), an emerging interoperability standard for connecting AI agents to enterprise systems. The company says the integrations also support developer environments Windsurf, VS Code, and Kiro, and can extend to any MCP-compatible platform.

The announcement reflects a widening concern in enterprise IT: agentic AI systems that draw on fragmented, ungoverned content repositories can produce contradictory outputs across workflows, creating compliance exposure that is difficult to reverse once agents have acted at scale. For organisations subject to regulatory recordkeeping requirements - in banking, government, legal, or healthcare - the risk is not hypothetical.

MCP has gained rapid adoption as an interoperability layer for AI agents. Anthropic – which developed the protocol – has confirmed adoption across platforms including Copilot, Cursor, Gemini, and Visual Studio Code.

eGain organises its AI Knowledge Connectors into four categories. Content Connectors draw from policy repositories, SharePoint, Confluence, CRM knowledge bases, and conversation archives. Data Connectors deliver contextual information to AI systems in real time.

Experience Connectors route verified answers to enterprise platforms including Salesforce, SAP, Zendesk, and contact centre environments such as Amazon Connect, Genesys, and Talkdesk.

Process Connectors apply identity controls, access rules, and business policies to ensure AI outputs remain within approved boundaries and generate an auditable confirmation trail.

The inclusion of Cursor and other AI-assisted developer environments extends the knowledge governance argument beyond service and operations teams. Governed enterprise knowledge can now inform how internal software is built, not only how it is used, according to eGain. That matters to enterprise architects and GRC managers because inconsistent knowledge feeding AI-assisted code generation can propagate errors across development pipelines with less visibility than errors in customer-facing AI.

https://www.egain.com

HiDock extracts Action Items from Conversations

Turning meeting transcripts into automated task lists is the promise behind HiNotes 3.0, a significant update from HiDock that repositions its note-taking platform from passive transcription toward post-meeting workflow management.

The company is making it available immediately to existing users at no additional charge. A Pro Membership tier offering advanced features is available, though pricing was not disclosed in the announcement.

HiNotes 3.0 introduces a redesigned three-pane interface consolidating notes, AI-generated tasks and schedules into a single view. The platform automatically extracts action items from meeting transcripts, organises them into prioritised task lists and maintains links to the original conversation context - addressing a persistent pain point for managers tracking decisions and commitments across distributed teams.

The update also introduces Smart Labels for automatic note categorisation, Speaker Suggestions that identify meeting participants based on prior meeting history and contact data, and Whisper Note Aggregation, which merges multiple standalone recordings into a single structured document. HiDock says these features reduce manual organisation effort.

HiNotes 3.0 supports seven AI models, including GPT-5.4, Claude Sonnet 4.6 and Gemini 3.1 Pro, according to the vendor. The ability to select AI models for specific tasks - summarisation, analysis or structured output - gives organisations a degree of control over their AI processing pipelines, a consideration increasingly relevant to enterprise AI governance frameworks.

The platform integrates with HiDock's P1 hardware device, which captures audio through personal earbuds anywhere.

HiNotes 3.0 is available immediately. The free tier includes transcription and AI summaries.

https://hidock.com

Cloud File Vendor Nasuni Retools for AI

Governing AI agent access to corporate file data has emerged as a pressing challenge for enterprise security and compliance teams. Boston-based Nasuni is attempting to address it with a new platform capability that lets AI tools operate directly on file data within existing permission structures - no separate data pipelines required.

The company, which provides cloud-native unstructured data management to more than 1,300 enterprises globally, has announced a brand refresh and two new product capabilities - AI Activate and Resilio Active Everywhere v6 - alongside a repositioning from cloud file storage vendor to what it describes as an "unstructured data platform for enterprise teams and AI."

Neither product is currently generally available. AI Activate is in invite-only preview, with general availability scheduled for Q4 2026. Active Everywhere v6 enters preview in Q2 2026, with general availability targeted for Q3 2026.

AI Activate uses Model Context Protocol (MCP) - an emerging open standard for connecting AI agents to enterprise tools and data sources - to extend governed file access to large language models and AI agents. Nasuni says AI tools can discover, read, and act on file data within existing trust boundaries, without creating data copies or separate pipelines.

The company claims AI agents would be subject to the same access controls as human users, reducing the risk of sensitive data being exposed to unauthorised AI queries.

Resilio Active Everywhere v6 - built on technology from Nasuni's acquisition of Resilio, - is designed to give distributed edge teams LAN-speed file access without WAN optimisation appliances or proprietary caching hardware.

The capability operates within Nasuni's existing namespace, permissions, and governance structure. Nasuni states this can reduce hardware costs for organisations with distributed workforces - a claim of particular relevance to IT infrastructure managers in sectors such as manufacturing, architecture, engineering and construction (AEC), and media, where large unstructured file sets are shared across multiple sites.

Founded as a cloud file storage company, Nasuni now describes itself as an unstructured data platform underpinning both human workflows and AI operations. CEO Sam King described the platform as "uniquely designed" to power enterprise AI .

https://www.nasuni.com/

DoxAI Fraud Check Lands in NZ

Artificial intelligence is being put to work on document fraud in New Zealand's financial sector, with FUJIFILM Business Innovation New Zealand (FBNZ) securing exclusive local distribution rights for Fraud Check AI, a product of Australian technology venture DoxAI.

Fraud Check AI automates the review of identity documents, payslips, financial statements, and supporting paperwork - scanning for inconsistencies, manipulations, and high-risk indicators that may signal fraud. The solution targets financial institutions handling large volumes of compliance documents where manual review is slow and error-prone.

The announcement comes as document-related fraud and identity crime continue to climb. New Zealand's Financial Markets Authority (FMA) reported that New Zealanders lost $265 million to fraud in the past year, with the government moving to introduce new anti-scam legislation under the Fair Trading Act following a parliamentary inquiry into bank fraud protections.

DoxAI, a venture of Sydney-based Lakeba Group, says Fraud Check AI has been trained on Australasian document samples and built to align with local regulatory and compliance frameworks. The company claims this regional focus delivers higher detection accuracy and faster deployment for ANZ financial institutions.

"Financial institutions in New Zealand are facing increasing pressure to detect and prevent fraud while handling thousands of documents every day," said Cameron Mount, FBNZ Director of Sales. "By partnering with DoxAI, we're expanding our capabilities and helping our customers reduce manual processes, strengthen compliance, and confidently identify fraudulent documents with AI-driven precision."

For compliance, records, and governance teams, the solution is positioned to reduce the administrative burden on staff conducting manual document checks during customer onboarding and lending processes.

FBNZ says the product supports compliance alignment to local regulatory requirements.

Beyond the financial sector, FBNZ indicates it plans to extend DoxAI's broader portfolio into legal, education, government, and healthcare settings - sectors where document-intensive workflows are common compliance pressure points.

"Our technology is built for the realities of the Australasian market," said Alfonso Porcelli, CEO of DoxAI. "By combining DoxAI's accuracy with FBNZ's process automation expertise, we're delivering a powerful solution that helps financial institutions detect fraud quickly and protect their customers."

Fraud Check AI is now available in New Zealand exclusively through FUJIFILM Business Innovation, delivered by the company's Process Automation team. DoxAI will continue to handle Australian and other international markets.

https://doxai.co/product/fraud-check-ai/

https://www.fujifilm.com/fbnz/en