




Active Metadata: The Key to Content Observability



Building an AI governance framework in seven phases

Why AI won't save your records program - and what will

The New Bottleneck for Enterprise AI Sits Inside the Document

The Best Cyber Security Find it before they do!

ezescan.
making digital work

PII/PCI Automated Discovery & Remediation

- ✓ Comply with data protection laws
- ✓ Reduce data breach risk
- ✓ Enhance customer/public trust
- ✓ Retrospective & real-time discovery



Directors Cannot Rely On Automation as a Legal Defence: Governance Institute

Boards need to treat the deployment of agentic AI as a delegation decision rather than a technology purchase, according to a new white paper from the Governance Institute of Australia. Unlike generative AI, which produces outputs for human review, agentic AI systems act on those outputs directly, transacting with customers and executing multi-step workflows with limited human oversight.

The white paper, [Governing in the Age of Agentic AI](#), was released by the Governance Institute of Australia with input from law firm Mallesons, SEEK, governance software firm Diligent, EthicAI and the University of Melbourne's Centre for AI and Digital Ethics.

It follows survey data cited in the paper showing just over half of organisations have already deployed AI agents.

The paper sets out five priorities for organisations deploying agentic systems: clear accountability, with an owner for every agent deployed; meaningful oversight through substantive human review at the systems level; defined authority boundaries on what agents can and cannot do; transparency and traceability of decisions that can be audited and explained; and capability uplift so governance literacy keeps pace with adoption.

Existing frameworks assume the wrong things

Daniel Popovski, senior policy and advocacy adviser at the Governance Institute, wrote in the organisation's [Governance Directions journal](#) that most existing governance frameworks assume predictable system behaviour, clear human oversight at decision points and bounded operational contexts.

"It is no longer sufficient to ask, is the output correct? We must now consider, was the decision authorised? Did it operate within approved boundaries? Can the organisation justify the outcome under scrutiny?" he wrote.

Popovski said governance must shift from model-level assurance, focused on accuracy, bias and explainability in individual systems, to system-level oversight recognising that risk arises from how agents interact across workflows, data sources and third parties.

He pointed to the risk of "automation bias", where humans defer to machine outputs as agents become more capable,



arguing this requires deliberate governance interventions including training, oversight tools and escalation mechanisms.

Agentic AI does not create a new legal entity, the paper notes, meaning organisations remain responsible for their AI systems' actions and directors retain their existing duties of care and diligence.

"The key takeaway is to move forward, but ensure governance keeps pace so AI can be deployed with confidence and control," said Bryony Evans, a partner at Mallesons.

Regulators flag cyber and legal exposure

The white paper notes that ASIC, APRA and the Australian Signals Directorate have warned that as AI becomes more autonomous, risk moves faster and becomes less predictable, requiring stronger cyber and resilience controls.

Agentic AI introduces new legal exposure, from privacy and data breaches to agents entering into contracts or breaching regulatory obligations, and liability for agentic outcomes remains untested in Australian courts.

The full white paper is available from the [Governance Institute of Australia](#).



Publisher/Editor: Bill Dawes

Email: idm@idm.net.au

Phone: 02 90432943

Web Development & Maintenance: Cordelta

Published by Transmit Media Pty Ltd

PO Box 392, Paddington NSW 2021, Australia

All material in Information & Data Manager is protected under the Commonwealth Copyright Act 1968. No material may be reproduced in part or whole in any manner whatsoever without the prior written consent of the Publisher and/or copyright holder. All reasonable efforts have been made to trace copyright holders. The Publisher/Editor bears no responsibility for lost or damaged material. The views expressed in Information & Data Manager are not those of the Editor. While every care has been taken in the compilation of editorial, no responsibility will be accepted by the Editor for omissions or mistakes within. The Publisher bears no responsibility for claims made, or for information provided by the advertiser.

Copilot overhaul after weak global adoption revealed

Microsoft is undertaking the biggest restructuring of its Copilot strategy since it began embedding OpenAI's technology across its software three years ago, a move triggered by [leaked internal figures](#) showing that very few of its commercial customers pay for the assistant or use it regularly.

The company will merge its separate consumer and enterprise Copilot products into a single unified application and introduce a new class of always-on agents called AutoPilots, while cutting features judged to have failed. Chief executive Satya Nadella confirmed the direction on Microsoft's 29 July earnings call, describing a forthcoming app where "chat, Cowork, code and Autopilots" come together for both consumer and commercial users.

The catalyst was an internal memo from Copilot chief [Jacob Andreou to his roughly 11,000-person team](#), first reported by in early July. The memo told staff the product must "earn the right to exist" by delivering measurable business outcomes, and set out figures that had not previously been disclosed publicly.

Fewer than 4.5 per cent of Microsoft 365's 450 million commercial seats have converted to a paid Copilot add-on, according to the memo, and only 20 to 30 per cent of those who do pay use it in a given week. Taken together, weekly active paid usage across the commercial base sits close to one per cent — a return the memo framed as unacceptable for a product embedded into every major Microsoft surface for nearly three years.

Microsoft's own research, cited in the memo, found that customers disliked switching between multiple separate Copilot products and could not identify clear value from the fragmented experience. The unified app is intended to address that complaint.

Under the overhaul, the split consumer and enterprise Copilot apps will be combined into [one product](#) spanning chat, Microsoft's Cowork experience, coding tools and the new agentic layer. Microsoft has not committed to a firm launch date, platform list or licensing model, saying only that the unified experience is due later this year.

As part of the reset, Microsoft is discontinuing features with low engagement. Andreou's memo named Copilot Podcasts, which auto-generated audio summaries of meetings and documents, and Copilot Labs, the channel used to test experimental features, as the first to be cut.

The restructuring comes despite growth in raw subscriber numbers. At its 29 July results Microsoft cited roughly 30 million paid Copilot seats and what it called its fastest three-month growth on record, indicating the problem is thin penetration and low engagement across a very large base rather than an absolute decline.

AutoPilots and autonomous action

The consolidation introduces AutoPilots, agents designed to run continuously in the background rather than respond to individual prompts. The first, [Microsoft Scout](#), was unveiled on 2 June and is in private preview.



Scout schedules and coordinates meetings across time zones, prepares briefing materials, tracks deadlines and flags risks, acting across Teams, Outlook, OneDrive and SharePoint under organisational controls.

The shift from prompt-and-respond assistance to always-on agents that read mail and calendars and act against SharePoint and OneDrive content raises questions that records and information management specialists have already begun to flag independently of Microsoft's plans.

[Gartner predicts](#) that by 2027, 40 per cent of enterprises will demote or decommission autonomous agents because of governance gaps identified only after production incidents, warning that failures are most likely when organisations "fail to distinguish between an agent's ability to act and the scope of access it is granted."

Senior director analyst Shiva Varma said the most autonomous agents - those that act within guardrails while humans review only exceptions and audit logs - demand the most rigorous controls, including continuous monitoring, enforced guardrails and clear ownership for agent behaviour.

In a July guide to governing Copilot agents, [AvePoint](#) set out the same fault lines in concrete terms: data exposure through over-scoped connectors, "compliance drift" when agents bypass the label, retention and data-loss-prevention policies that apply to the underlying content, and "identity ambiguity" when an agent acts without a clearly accountable human owner.

Writing for [AIIM](#), the professional body for information management, Shell information-management lead Jean-Paul Meijer similarly warned that an agent "operating inside a governance vacuum" creates risk rather than value. Microsoft has said AutoPilots operate under organisational controls but has not detailed enterprise-management specifics.

POWERED BY
ezescan.



Smarter. Faster. Easier. Records Capture & Process Automation

- ✓ AI Assisted Document Classification
- ✓ Seamless EDRMs Integrations
- ✓ Automated Email / eForms Capture
- ✓ Digital Mailroom Automation
- ✓ Simplified Back Scanning

Call: 1300 EZESCAN (1300 393 722)

www.ezescan.com.au

Data Breach Notifications Hit Record Highs in A/NZ

Data breach notifications hit record highs on both sides of the Tasman in 2025, with regulators in Australia and New Zealand both pointing to cyber attacks and vendor security gaps as the common thread.

Australia recorded 1,205 notifications under the Notifiable Data Breaches (NDB) scheme in the 2025 calendar year, the highest since the scheme began in 2018 and an 8 per cent rise on 2024's 1,112. The figures were published by the [Office of the Australian Information Commissioner \(OAIC\)](#) on 6 July 2026, drawing on notifications from businesses and Commonwealth agencies covered by the Privacy Act.

Underlying [NDB scheme data](#) for July to December 2025 shows 670 notifications in that six-month period, with malicious or criminal attacks accounting for 405 of them, more than double the 194 attributed to human error.

Health service providers were the most commonly affected sector across the year, accounting for 225 notifications, or 19 per cent of the total. Financial services (157 notifications), the Australian Government (118), business and professional associations (103), and education and legal, accounting and management services (81 each) rounded out the top five sectors by volume.

Health data was also at the centre of New Zealand's biggest privacy failure of the year. A December 2025 cyber breach of the Manage My Health patient portal compromised the health records, clinical notes and identity documents of close to 100,000 New Zealanders, with the stolen data later offered for sale on the dark web. It is New Zealand's largest-ever health data breach.

Cyber hacking remains the leading cause

Across the full year, malicious or criminal activity was attributable for 716 of the 1,205 Australian notifications. Within the July-to-December data, cyber incidents alone accounted for 253 notifications, ahead of social engineering or impersonation attacks (88) and rogue employee or insider threats (34).

"The threat posed to Australian businesses and organisations by data breaches is substantial and rising year on year, with 2025 recording the highest number of notifications received in a year since the commencement of the NDB scheme," said Australian Privacy Commissioner Carly Kind.

The OAIC's 2026 Australian Community Attitudes to Privacy Survey found data breaches are now the top perceived privacy risk among Australians, with 82 per cent concerned about the issue, up from 74 per cent in 2023.

New quick reference guide for entities

Alongside the statistics, the OAIC published a new [quick reference guide](#) for entities with NDB scheme obligations. The guide sets out four steps for responding to a breach: contain it to prevent further unauthorised access or disclosure; assess the incident to determine whether it is an eligible data breach; notify affected individuals and the OAIC where required; and review the incident to reduce the likelihood of a repeat.



It includes a decision-making framework and self-assessment checklist to help entities determine whether a breach is likely to result in serious harm and therefore requires notification.

"With this new guide, entities subject to the Privacy Act will have quick access to essential information needed to act effectively when faced with a potential data breach, when they may be under significant pressure," Commissioner Kind said. "This will benefit both the entity, and the impacted community."

The guide is available as both an interactive webpage and a downloadable [self-assessment checklist](#) (PDF, 261 KB) from the OAIC website.

New Zealand sees breach notifications climb 27pc

New Zealand's Office of the Privacy Commissioner (OPC) reported a similar trend in its Annual Report 2024/25. The office received 1,093 privacy breach notifications in the year to 30 June 2025, up 27 per cent on the prior year, and 591 of those were classified as serious privacy breaches, a 43 per cent increase. Total privacy complaints from individuals rose 21 per cent to 1,598, a record high for the office.

"This year we also had a 27% increase in privacy breach notifications. It's clear that agencies need to respond to the challenge to be better at safeguarding New Zealanders' personal information," Privacy Commissioner Michael Webster said in the report.

A March 2025 survey of more than 1,000 New Zealanders found 82 per cent wanted more control and choice over the collection and use of their personal information, mirroring the concern levels reported in Australia.

Unlike Australia's NDB scheme, New Zealand's Privacy Act 2020 carries no financial penalty regime for breaches, a gap the Commissioner has repeatedly flagged as reducing agencies' incentive to invest in privacy safeguards.

AI Risks Could Turn Catastrophic by 2030

There is a very real risk the AI revolution could cause an international catastrophe within five years, a panel of 272 experts has found. The experts defined catastrophe as more than one million deaths, more than \$US100 billion in losses, or civilisation-scale damage to democracy or privacy.

MIT FutureTech and the MIT AI Risk Initiative ran the study, published in June 2026. The panel drew AI researchers, policy advisers, technologists and governance specialists from 37 countries. They came from academia, industry, government and civil society.

The experts were presented with 24 risk categories, from discrimination and privacy loss to weapons, cyberattacks and power centralisation. They judged 18 of these more than 10 per cent likely to cause catastrophic harm by 2030.

The panel mixed global heavyweights with Australian institutions. Oxford, MIT, Stanford, Harvard and Princeton were all represented. Three of the study's core authors are based at the University of Queensland. The University of Sydney and the University of the Sunshine Coast also contributed experts. Two staff from Commonwealth Bank of Australia joined the panel.

The researchers used the Delphi method, a technique built for forecasting in fields with little hard data. Experts rate each risk anonymously across several rounds, then see how peers reasoned and revise their answers. The method surfaces genuine disagreement and limits pressure to conform. Here it ran three rounds in late 2025.

Under many risk-governance frameworks, a 10 per cent chance of catastrophe over five years is deemed "intolerable". Such a level would normally trigger mandatory mitigation.

The findings sit within a wider debate about AI's extreme tail risks. An earlier tournament by the Forecasting Research Institute put the chance of human extinction from AI this century near 3 per cent.

Asked to name their top three concerns across all 24 risks, the experts scattered rather than agreed. Weapons and cyberattacks topped the list at 26.8 per cent, ahead of power centralisation at 23.5 per cent. Disinformation and influence reached 22.1 per cent. Loss of consensus reality and dangerous capabilities each drew 21.6 per cent.

One of the experts commented, "A technology-loving researcher develops AI with only technology in mind. He completely ignores the transparency and explainability of the model and pushes ahead with the development of multi-agent systems.

"If multi-agent systems were distributed across the world, who would be able to audit them? Suppose agents were distributed across the United States, Europe, and Asia, and they were linked together. If 100 or 200 agents were linked together and operating, no one would know what was going on."

Severity was scored from one, negligible, to five, catastrophic. Experts spread probability across both the catastrophic tail and the middle range. That pattern signalled doubt about the scale of harm, not whether it would occur. The highest-severity risks were dangerous capabilities, weapons and cyberattacks, power centralisation, competitive dynamics and false information.

Modelling cost-effective mitigations lowered every risk, but five stayed above 10 per cent. All 24 stayed above 5

per cent. The stubborn five were dangerous capabilities, weapons and cyberattacks, environmental harm, inequality and unemployment, and power centralisation.

Across 14 sectors, experts rated information and national security as the most exposed. Finance and insurance followed closely. The most exposed sectors were those where AI is embedded in critical decision-making, the report says. They named industries "where failures have immediate, large-scale consequences".

Finance and insurance drew high marks for fraud, security flaws and system failures. A Delphi participant cited "both direct attack vectors (fraud, market manipulation) and regulatory exposure from AI system failures".

The full report is available at <https://airisk.mit.edu/priorities>.

Food Sector Cyber Risk Exposed By Sugar Hit

A ransomware attack on Australia's second-largest sugar producer shut two Queensland mills, the latest blow to critical infrastructure as attacks on essential industries climb.

Mackay Sugar confirmed a cyber security incident affecting parts of its operations. Milling stopped at the Farleigh and Racecourse plants near Mackay during peak crushing season. The strike fits a widening pattern. Ransomware crews increasingly target food, energy and utilities, where downtime is costly and operators face pressure to pay.

The [Food and Ag-ISAC](#) recorded 265 ransomware attacks on the sector in 2025. Total incidents across all sectors rose 82 per cent to 6,377.

A third Mackay Sugar mill at Marian was unaffected because it had not yet begun crushing. The company runs three mills and produces about 700,000 tonnes of raw sugar each year. The attack forced roughly 1,300 mostly family-owned farms to stop harvesting cane. Cut cane loses sugar content quickly, so the halt threatened grower income across the region.

Mackay Sugar told growers to cease harvesting immediately and wait for further instruction. Industry body Canegrowers Mackay relayed the directive to its members.

The producer engaged specialist cyber security experts and notified authorities. It said its priorities were staff safety, protecting operational systems and maintaining business continuity. Mackay Sugar later reported significant progress restoring cane supply, harvesting and milling systems. Steam trials began, with a staged crushing restart expected within days.

A ransomware group calling itself The Gentlemen named Mackay Sugar on its Tor-based leak site. It set a countdown timer but had not published any data. Microsoft tracks the group as Storm-2697. Researchers describe an affiliate model, double extortion and malware that moves rapidly across networks.

It remained unclear whether attackers reached industrial control systems directly, or whether disruption followed from compromised IT systems.

Critical infrastructure now draws sustained attention from criminals, hackers and state-backed actors seeking disruption and leverage.

Researchers have identified 72 active threat actors targeting food supply chains globally. Russian-linked groups account for most observed activity, followed by Chinese-linked groups.

How One Phone Call Exposed 5.7 Million Qantas Records

Qantas has dodged penalties that could have run into billions of dollars after simple human error was found to be the cause of the massive 2025 data breach that affected millions of customers, and ISO27001 compliance failed to prevent.

The Office of the Australian Information Commissioner (OAIC) closed its preliminary inquiries on 16 July 2026 without launching an investigation or taking any regulatory action, finding the airline could not reasonably have foreseen or prevented the attack.

The OAIC report gives the most detailed official account yet of the incident, which struck the Salesforce customer relationship management (CRM) platform used by Qantas contact centres.

On Saturday 28 June 2025, a threat actor phoned an agent at an overseas contact centre run by a third-party provider for Qantas. The caller posed as ‘Qantas IT help’ checking an IT issue the agent might have been experiencing. The agent believed the caller was a legitimate Qantas employee.

The caller directed the agent to a website related to the Salesforce platform and walked them through a series of actions supposedly needed to close an IT support ticket. In the wider campaign this attack belonged to, victims were directed to the Salesforce connected app setup page and told to enter an eight-digit connection code.

That code registered an attacker-controlled tool, a modified version of Salesforce’s own Data Loader utility, as a trusted connected app on the victim’s account.

The result was that the agent’s own authenticated session authorised the connection. No password was stolen. No multi-factor authentication challenge was triggered, because the attacker never logged in. Salesforce saw only a legitimate user approving an app, and issued the attacker an OAuth token carrying all of the agent’s permissions.

Those permissions were broad. As a customer service agent, the employee was authorised to view the contact profiles of Qantas customers across the CRM. The attacker’s tool used that access to run bulk queries through the Salesforce API, extracting approximately 5.67 million customer records, including those of about 5.12 million Australians.

For around 4 million records the data comprised names, phone numbers, email addresses and Frequent Flyer numbers, tiers, points balances and status credits.

A further 1.7 million records also held one or more of residential or business addresses, dates of birth, gender and meal preferences. No credit card, financial or passport details were stored on the platform, and no passwords or PINs were accessed.

The breach occurred despite Qantas requiring its overseas contact centre provider to maintain ISO 27001 compliance, a factor the Office of the Australian Information Commissioner (OAIC) described as “particularly relevant” in assessing whether the airline took reasonable steps.



The regulator noted the standard mandates controls directly connected to the incident, covering security awareness training on phishing and social engineering, rules for managing access credentials, and data leakage prevention monitoring.

The intrusion generated two system alerts from an unusual number of failed login attempts, incidental noise rather than the access vector itself. Because the attack landed on a Saturday, the alerts sat unread until Monday morning, 30 June, when a Qantas staff member spotted them and escalated to the cybersecurity team.

Qantas froze and revoked the compromised account the same day, assessed the data exfiltration, and triggered its incident response. It disclosed the breach publicly on 2 July and notified all affected customers of the specific data types taken from 9 July.

The attack was part of a global vishing campaign against Salesforce customers attributed to the ShinyHunters group, which also claimed victims including Google, Allianz Life and LVMH brands.

The Qantas data was published by the hybrid Scattered Lapsus\$ Hunters extortion group on 12 October 2025 after a ransom deadline passed, despite a NSW Supreme Court injunction prohibiting its release.

Root Causes

The OAIC identified the central vulnerability as a default Salesforce configuration that allowed any end user to authorise the connection of a third-party application. Role-based access controls could not have prevented the breach, because the attacker’s tool simply inherited the legitimate access of a deceived user. Salesforce has since changed that default for all its customers.

The second root cause was a training blind spot the

regulator said extends across the industry.

“Most social engineering training focusses on credential theft and does not usually address the less common tactic of inducing an employee to authorise access through legitimate system interactions,” the report states. The attack would likely have succeeded even with standard training in place.

Two further factors amplified the damage. A single contact centre agent could see the entire customer base, so one deceived employee exposed 5.7 million records. And detection relied on login-based alerts reviewed on the next business day, giving the attacker an undisturbed window of roughly 36 hours over a weekend.

What Could Have Prevented It

The single control that would have stopped the attack outright is the one Salesforce has now made default: restricting connected app authorisation to administrators, so end users cannot approve third-party API connections. Organisations running Salesforce or similar CRM platforms can go further by allow-listing approved API clients and limiting the API Enabled permission to roles that need it.

Realtime exfiltration monitoring would have cut the breach short. Bulk extraction of millions of records via API is exactly the anomaly that Salesforce Event Monitoring and transaction security policies can flag, with automated session termination, rather than waiting for a human to read login alerts on Monday. Around-the-clock alert triage matters for the same reason, as attackers deliberately strike on weekends.

Procedural controls also had a role. A mandatory callback protocol for unsolicited IT support calls, requiring agents to verify a ticket number in the service desk system or ring back a known internal number, directly counters the pretext used. Rate limits or approval gates on mass record retrieval would have shrunk the blast radius even after the authorisation succeeded.

Since the incident Qantas has rolled out a dedicated online social engineering course for contact centre agents, reinforced with face-to-face training, addressing the consent phishing gap the OAIC identified.

The airline had spent a year exposed to the toughest penalty regime in Australian privacy law. Under section 13G of the Privacy Act, a serious interference with privacy carries a maximum penalty of the greater of \$A50 million, three times any benefit obtained, or 30 per cent of adjusted turnover where that benefit cannot be valued.

As IDM [reported in July 2025](#), the turnover limb implied a theoretical maximum of about \$A6.6 billion on Qantas’ FY24 revenue, and more than \$A7.1 billion on FY25 revenue of \$A23.8 billion.

The realistic ceiling was likely \$A50 million per contravention, since Qantas obtained no benefit from the attack, but contraventions can multiply. In its ongoing Medibank proceedings, the OAIC has treated each affected individual as a potential separate contravention. In the end, Qantas will pay nothing.

The regulator stressed that its report makes no concluded findings and is not an endorsement of Qantas’ practices. It remains open to the OAIC to investigate later, and a representative complaint and individual complaints over the breach are still being handled separately.

“Based on the information provided, it does not appear that Qantas could have reasonably foreseen and prevented the breach in the manner that it occurred,” the Commissioner concluded.

The report is available at [oaic.gov.au](https://www.oaic.gov.au).

Parliament House Falls Short on Essential 8

The department that operates the Australian Parliament House’s computing network has implemented 7 of the mandatory Essential 8 cyber security strategies below the required standard, the Australian National Audit Office has found.

The audit assessed whether the Department of Parliamentary Services (DPS) has an effective baseline of cyber security controls. The ANAO concluded the baseline was only partly effective.

DPS provides IT services to 4,817 users and manages 10,951 end-user devices across the Parliamentary Computing Network. Its users include parliamentarians, their staff and other parliamentary departments, each with differing business and security needs. The department was the subject of a high-profile cyber security incident in 2019.

Under the Protective Security Policy Framework (PSPF), non-corporate Commonwealth entities must implement the Essential Eight strategies to Maturity Level Two. The ANAO found DPS was relying on risk-management approaches to cover the gap between implemented and required controls. For the seven strategies using risk mitigations, those approaches fell short of the standard needed to address the risk.

“The Essential Eight cyber security strategies have not been fully implemented in accordance with the requirements of the Protective Security Policy Framework, primarily because the department is relying on compensating controls without adequate coverage of all systems and risk management of identified vulnerabilities,” the report states.

The audit found governance weaknesses familiar to information management professionals. DPS had not completed identification and documentation of key systems and ICT assets. Inventories of key information, assets and risks were outdated. Key cyber security policies were incomplete or not endorsed by the department’s Security Committee.

The department lacked a single source of truth for identified risks and issues, and available registers were not complete. Limited evidence supported previous PSPF self-assessments, although the most recent assessment in 2024-25 was backed by control testing. The Information Security Working Group, the body meant to coordinate cyber security and business alignment, had not met in the 12-month period examined. The ANAO also found DPS accepted risks above its own tolerance when approving new technology systems. Systems in active use required re-assessment and approval. Significant ICT staff turnover over the previous 18 months complicated ongoing risk management.

The ANAO made two recommendations. DPS should review its governance arrangements and risk assessment processes, and develop and implement a prioritised, risk-based program of uplift activities to achieve PSPF compliance. DPS agreed to both.

Why AI won't save your records program - and what will

By Nigel Carruthers-Taylor

In a widely-shared ARMA webinar, Kaan Volkan made this clear warning: AI large language models, applied without adequate information architecture, fall dangerously short for records programs - and what that can cost organisations that fail to heed the warning, including falsely destroyed records. His presentation title said it plainly: “LLMs Destroyed My Records.”

I am not surprised. Forty years in this work, and I've watched this pattern repeat too many times to count - but never at this speed, and never this final. The stakes have never been higher. These records could be wrongfully destroyed, and they are not merely operational assets. Often they are the evidence trail of decisions made by public institutions - institutions accountable to the people they serve. Destroy those records, and you destroy the ability to hold power to account. That is not a records management problem. It is a democratic one.

This is the argument I am making in my essay, Records, Accountability and Democracy. My central contention is this: records management is not an administrative discipline. It is the cornerstone of democracy. Right now, we are failing it.

The frictionless trap

Think about what democratic accountability actually requires. A citizen exercises their right to information. A parliamentary committee calls for the papers. A royal commission follows the thread of a decision back to its origins. An auditor reconstructs how a contract was awarded. Every one of these processes depends - entirely and without exception - on records having been made, kept, and preserved in a form that is authentic, reliable, complete and accessible. When records fail, accountability does not merely become difficult. It becomes impossible.

The paradox of our era: we have never created more information, and never had a harder time finding it. I call this the frictionless trap. When creating a document costs nothing, when saving or sharing a file takes one click, and when every message is automatically retained somewhere, then the natural response is to create more, save more, share more, and organise nothing. A single procurement decision may now be scattered across Teams messages, Outlook threads, SharePoint documents, OneDrive folders, Salesforce records, SAP entries, and mobile phone photos - with no authoritative record of what was decided or why.

We have seen this play out with disturbing regularity: ministers who cannot produce documents showing how decisions were made, procurement processes whose evidentiary trail has been scattered across personal email accounts and encrypted messaging apps, public health deliberations that are simply unrecoverable. These are not administrative failures. They are democratic failures - and they are being produced, in large part, by the collapse of records management discipline in the digital age.

Compounding the technical fragmentation is a

“A government that claims its records are managed by AI - but whose information architecture has never been analysed - has not improved its accountability. It has substituted the appearance of control for the reality of it.”

cultural one that I think we underestimate. Staff in contemporary digital workplaces rarely think about evidentiary completeness - they think about getting work done. As AI is increasingly involved in decision-making, this gap will widen. Records that would explain how an AI-assisted decision was reached may simply never be created, leaving organisations unable to account for choices that may have profound consequences for the people they serve.

Architecture first - always

The profession's response to fragmentation has been AI autotclassification and manage-in-place governance. Both are genuine advances, and I am not dismissing them. However, when adopted as complete solutions rather than carefully bounded components of a broader strategy, they can make the democratic accountability problem significantly worse than the one they set out to solve.

Research from Cornell University found that AI systems hallucinate at rates of 69 to 88 percent on complex tasks, and up to 50 percent on relatively simple ones. In a records context, this is not mere operational inefficiency. An AI that mis-classifies a document - applying a short-term retention schedule to a high-value record, or failing to recognise that a cluster of emails constitutes the evidential record of a significant decision - is potentially destroying the very evidence that democratic accountability will one day demand. The AI does not know it is dismantling an audit trail. But that is precisely what it does.

The root cause is almost always the same: AI was applied before the information architecture was in place. I experienced this directly while developing an architecture for a government agency. I persuaded them to build information categories around how people actually worked, rather than using inherited records classification schemes that staff neither recognised nor used.

AI could then be trained to work with the grain of the business - and, critically, reflect the actual shape of the decisions the organisation was making and needed to be held accountable for. The organisation's IT team subsequently declared the approach 'too complex and old school' and instructed their AI to classify existing unstructured content instead. Very little progress followed.

The same failure mode applies to manage-in-place governance more broadly. Tools like Microsoft Purview offer real advantages; they scale, reduce friction, and



integrate with the platforms that people actually use. On the other hand, they were not designed to maintain the contextual integrity that high-value, accountability-significant records require.

The context of the aggregation of related items across systems, defensible proof of disposal, and the full evidentiary record of a complex decision. A global IG World survey found Purview adoption for governance had actually declined - from 50 to 40 percent of organisations - with 94 percent reporting significant challenges.

That is not a failure of education. It is practitioners arriving, through experience, at the same conclusion: manage-in-place, taken on its own, is not adequate for the kind of defensible record-keeping that public accountability demands.

“An un-auditable system of government is not a democracy in any meaningful sense. It is a managed fiction.”

The hybrid answer

My answer is not the centralised EDRMS model or outsourcing records management to applications. Nor is it pure manage-in-place. It is a deliberate hybrid: the right model for the right record, matched to the accountability that record carries. Years of implementation and reading brought me here - not theory alone.

High-value records - those that document decisions, transactions, and the exercise of power - are centralised in a purpose-built EDRMS with defensible lifecycle management. Repeating transactional processes can embed records management at the point of creation. Routine, low-risk content is governed in place, with the limitations of in-place tools understood and mitigated.

In all cases, information architecture must come first. Not as preliminary work before the records management strategy but as the core strategy itself. Let me be emphatic about that.

Get it right, and AI classification becomes a powerful, scalable tool: governance applied at the volume modern organisations demand. Get it wrong, or skip it, and AI embeds the records problem at speed and scale, under a false veneer of automated compliance.

The accountability gap doesn't close. It widens, invisibly, behind a dashboard that says everything is in order.

Why I'm writing this now

I am reflecting on this at a moment when the world is producing governments and institutions that are increasingly comfortable with opacity. They are content to operate in environments where the records of how decisions were made is inaccessible, incomplete, or simply non-existent. AI will accelerate that tendency if we let it. It will not, on its own, reverse it.

In my full essay, I set out six principles for practitioners and make the case, with the conviction of forty years, that records management is one of the most important things that a democratic society can invest in.

It isn't glamorous and it rarely gets its due. Records managers are the custodians of institutional memory. Institutional memory is the currency of accountability. Without it, the promise that power can be seen, examined, and held to account is nothing but mere words.

Records, Accountability and Democracy: Reflections After Forty Years is published in full on [my blog](#).

Nigel Carruthers-Taylor is the retired founder and Principal of iCognition, acquired by Kapish in April 2026. He is the author of a series of practitioner articles on hybrid records management, AI in records practice, and the use of Microsoft Purview and EDRMS in the Microsoft 365 environment, published in IDM and elsewhere.

The New Bottleneck for Enterprise AI Sits Inside the Document

By Sid Manchkanti and Abhas Ricky

For most of the last two years, enterprise AI conversations started with the model. Organisations debated which foundation model to use, how to fine-tune it, and which orchestration framework would deliver the best results.

That conversation is changing. As foundation models become more capable, accessible, and interchangeable across providers, many organisations are discovering that model performance is no longer the primary constraint on AI outcomes. Instead, the bottleneck has moved earlier in the pipeline: into the document layer that feeds AI systems in the first place.

In conversations with CIOs and CDOs across financial services, healthcare, and telecommunications, the same observation comes up repeatedly. The challenge is no longer how the model reasons. It’s what the model is reasoning.

The model is no longer the bottleneck for enterprise AI inference. The document understanding layer is. That sentiment reflects a structural shift in where value and risk now sit in the enterprise AI stack.

The Real Problem: Unstructured Data and Document Intelligence

In regulated enterprise environments, most critical data does not live in clean, structured warehouse tables. It lives in unstructured formats: PDFs, scanned filings, claims schedules, contracts, financial statements, lab reports, and rate exhibits. This is the data that feeds AI systems.

Document intelligence refers to the process of converting this unstructured data into structured, usable inputs for AI models. When this step fails, the consequences ripple through the entire AI pipeline.

The failure mode is often deceptively simple. A misread table or merged cell creates a malformed extraction. That extraction produces flawed embeddings, which return the wrong context during retrieval. The model then generates a confident answer from a confidently wrong input.

At that point, even the most advanced model cannot compensate for the underlying error. Improvements to model performance do not correct a structural issue that occurred before the model ever ran. In practice, the quality of the document pipeline often determines the AI system’s accuracy ceiling.

When Parsing Errors Become Business Risks

The impact of poor document parsing shows up directly in business outcomes. They are concrete, measurable, and largely underappreciated at the executive level. This is what makes document intelligence more than a technical challenge. For many organisations, it has become an operational and financial one.

Parsing errors are rarely visible when they occur. Instead, they compound downstream through workflows, decisions, and business processes. By the time the issue surfaces, the cost of remediation is often far greater than the cost of prevention.



In financial services, a single misread value in a fund administrator’s capital account statement can cascade into downstream errors in underwriting or reserving models. These errors can carry regulatory implications and lead to costly remediation efforts that often run into the millions.

Healthcare organisations continue to rely heavily on manual document abstraction for claims, remittance advice, and clinical documentation. This is driven in part by the structural complexity of the data, and in part by strict requirements around protected health information.

Manual document abstraction is consistently one of the largest line items in health data operations budgets.

In telecom, vendor interconnect billing and service-level agreements (SLAs) often contain complex rate tables that few systems read accurately at scale. Even small inaccuracies can translate into hundreds of millions of dollars of leakage at carrier scale.

The pattern across these industries is the same. Inaccurate document understanding is not a technical inconvenience. It is a P&L problem that quietly compounds.

The Trade-Off: Accuracy vs. Data Sovereignty

Layered on top of the accuracy problem is a second constraint specific to regulated enterprises: where AI processing happens.

Over the past several years, most of the enterprise AI inference stack has steadily moved into controlled customer environments – virtual private clouds (VPCs), on-premise infrastructure, or sovereign cloud regions. Models, vector stores, orchestration layers, and observability now operate within the same governance controls as the underlying data. Document parsing has been the forced exception.

Historically, the most accurate document processing options were delivered via SaaS APIs only, which left regulated customers choosing between accuracy and sovereignty:

■ Route the most sensitive documents in the enterprise out to a third-party API for higher accuracy, or

■ Keep data within the enterprise boundary and accept a meaningful accuracy gap on the workflows that matter most.

Compliance, legal, and risk teams have long viewed both options as compromises. As a result, many organisations have struggled to balance two equally important priorities: achieving the accuracy required for business-critical workflows while maintaining control over where sensitive data is processed.

Until recently, there was no clear path to achieving both.

A Maturing Category for Enterprise Document Intelligence

The good news is that this trade-off is beginning to close. Across the industry, organisations are applying greater rigour to how document intelligence systems are evaluated, particularly for complex tables and highly structured business documents that have historically challenged traditional parsing approaches.

At the same time, a new generation of document intelligence providers is making it possible to achieve high levels of parsing accuracy within customer-controlled environments. Recently, the team at Pulse open-sourced PulseBench-Tab, a frontier benchmark for table parsing built specifically around the kinds of documents regulated enterprises actually run on.

It contains 1,820 human-annotated tables drawn from real financial filings, government reports, corporate disclosures, and regulatory filings, spanning 9 languages and 4 scripts, many of which contain merged or spanning cells and complex structures that commonly break traditional parsing systems.

Importantly, the benchmark introduces T-LAG, a unified scoring approach that captures both text and structural accuracy. This ensures that systems are not rewarded for extracting approximate text while silently breaking the table’s shape.

Results from this benchmark show that frontier-level accuracy in document parsing is now achievable without a third-party SaaS endpoint, bringing a new level of reliability to enterprise AI pipelines.

Nine providers were evaluated independently and in the open, and the methodology benefited from academic contributions from members of S&P Global’s Enterprise Data Organisation. On that benchmark, Pulse delivered a T-LAG score of 0.9347 with full coverage across all 1,820 samples, materially ahead of the next closest provider at 0.8155.

Bringing Document Intelligence Inside the Enterprise AI Stack

This progress unlocks a new architecture for enterprise AI – one where document intelligence operates within the same environment as the rest of the data pipeline. As document intelligence becomes deployable within governed enterprise environments, organisations gain the ability to bring document processing into the same operational and governance boundary as the rest of the AI stack.

Combined with an AI-powered lakehouse architecture, this creates a more unified approach to managing structured and unstructured data, with consistent security, lineage, observability, and governance controls from ingestion through inference.

Solutions such as Pulse demonstrate what this architecture can look like in practice, enabling organisations to parse and structure complex documents without requiring sensitive data to leave the enterprise environment.

The result is a fully integrated pipeline that can:

- Parse unstructured documents
- Convert them into structured data
- Embed and retrieve relevant context
- Generate outputs using AI models

All within the same controlled environment.

For the CIO, that means a single governance boundary across the AI workflow rather than a patchwork of disconnected environments to secure, audit, and manage.

For the CFO, it can shift document processing from a recurring external service cost to an internal capability built on infrastructure that already supports broader AI and data initiatives. More importantly, it changes where organisations should focus their investments. As models become increasingly accessible, competitive advantage is shifting toward the quality, governance, and reliability of the data pipeline that powers them.

What This Unlocks for Regulated Industries

For executives setting AI strategy in regulated industries, improvements in document intelligence create a visible impact on operating metrics.

Financial services teams can keep 10-K analysis and filings, fund administration records, bordereaux processing, claims schedules, and actuarial reports entirely within their governed environment, with structural accuracy high enough that downstream agents can be trusted with the output, reducing pressure and time spent on human review cycles. Headcount that was previously dedicated to manual reconciliation can be redirected to higher-value analytical work.

Healthcare organisations can automate document-heavy workflows like clinical trial data extraction, lab panel ingestion, and explanation of benefits (EOB) processing into the same environment as their structured PHI. This materially reduces one of the largest line items in health data operations while accelerating revenue cycle times and clinical research workflows.

Telecom operators gain the ability to accurately interpret interconnect agreements and billing structures at the level of detail required to recover the revenue leakage that has historically been buried inside complex rate tables.

In each case, improved document intelligence directly translates into measurable business value.

The Future: The Sovereign AI Stack

The centre of gravity in enterprise AI is shifting. As models continue to converge in capability, the durable competitive advantage is moving one layer down, into the data to inference pipeline – specifically, how effectively organisations can process and govern unstructured data.

Document intelligence now sets the ceiling for accuracy and ROI. At the same time, data sovereignty is non-negotiable in regulated industries: AI must run where the data lives. This is where Cloudera’s AI and data anywhere vision applies: deploy AI across hybrid and multi-cloud environments, keep data in place, and enforce consistent governance.

Combined with Pulse, the regulated enterprise has a path to AI Native that protects accuracy, control, and the underlying ROI of every workflow built on top. That is the sovereign AI stack our customers have been asking for, and it is now within reach.

Sid Manchkanti is Co-founder and CEO, Abhas Ricky is Chief Strategy Officer at [Pulse AI](#). Originally published [here](#).

Most Unstructured Data Unready for AI

Seventy percent of organisations say less than half their unstructured data is discoverable and usable for analytics or AI, according to new research from analyst firm BARC. The finding lands as enterprises push AI agents into production on top of documents, emails and images they cannot fully see. The study, Harnessing Unstructured Data for AI Innovation, surveyed 225 data, analytics and AI leaders across North America and Europe in February and March 2026.

The research exposes a sharp gap between confidence and capability. Four out of five respondents are confident they can extract value from unstructured data without compromising governance controls. Yet roughly one third report inconsistent or non-existent controls for data bias and data lineage. About one third cannot consistently enforce data access and usage policies.

“More troubling still, some AI adopters still ‘don’t know what they don’t know,’” the report states. “Automated decisions and actions by agents only compound the potential damage of poorly governed unstructured data.”

The stakes are rising as autonomous agents act on data without a human in every decision loop. The authors warn that laggards “must pay close attention and move fast to catch up, because unstructured data provides the deep context that AI agents need to take smart, safe actions.”

The findings echo a recent [Cloud Security Alliance survey](#), which found more than half of organisations have only partial visibility into where unstructured data is stored. Despite the gaps, BARC found surprising progress on governance. Across every category, more respondents have implemented or optimised governance controls for unstructured data than a year earlier. Data quality controls jumped to 47 percent implemented or optimised, up from 36 percent last year.

Maturity remains uneven. North American respondents outpace Europeans in every governance category, including regulatory compliance at 58 percent implemented or optimised versus 48 percent. “This begs the question of whether General Data Protection Regulation (GDPR) and the EU AI Act have achieved their privacy and safety objectives,” the report says.

Asked how they measure the AI readiness of unstructured data, respondents pointed to accuracy (54 percent), consistency (48 percent), business relevance (44 percent), completeness (39 percent) and timeliness (31 percent). The report notes these measures require probabilistic logic and judgement rather than deterministic quality rules. That contrast with structured data keeps humans in the mix, with human validation ranking high at 41 percent.

Skills and training top the list of obstacles, cited by 36 percent of respondents, ahead of data privacy (30 percent) and data quality (29 percent). Data quality also dominates success measures at 48 percent, as organisations recognise the risks of faulty model inputs.

Classification (60 percent) and validation (56 percent) are the most widely adopted tasks for preparing unstructured data, well ahead of formatting, cataloguing and indexing. Vectorisation trails at just 17 percent. “You need your map before you begin the voyage,” the authors write.

Generative AI dominates model deployment, with 73 percent of respondents using generative models with unstructured data. Microsoft Azure leads platform choices at 52 percent, ahead of AWS (38 percent) and Google Cloud (32 percent). More than half of databases containing unstructured data reside on-premises or in hybrid environments.

Given that distribution, BARC recommends organisations build toward a platform-independent semantic layer covering federated querying, aggregated data views and integrated metadata. The firm expects adopters to favour platform-independent tools that support open ecosystems and eliminate lock-in risk.

The full report is available [here](#)

AI Vendor Lock-In Now a Risk

Most enterprises are locked into AI systems they cannot easily change, and the overwhelming majority do not fully understand their own AI dependencies, a new global study warns. The findings elevate AI sovereignty from a technical concern to a business continuity issue.

The Calculus of AI Sovereignty study, from the IBM Institute for Business Value in collaboration with Oxford Economics, surveyed 1,000 senior executives across 16 countries and 17 industries between February and April 2026. Respondents were responsible for AI, data, technology or related enterprise capabilities. Some 71 per cent of respondents say switching their primary AI vendor or model would be difficult. A further 68 per cent say meeting data residency and sovereignty requirements across geographies is challenging.

Visibility is the sharper problem. Ninety-one per cent say they do not fully understand their organisation’s dependencies across AI vendors, models and infrastructure, limiting their ability to assess risk and plan for disruption. Leaders report an average of six AI-related disruptions over the past two years, largely driven by vendor services. Even so, 81 per cent say a seven-day vendor outage would still cause severe or critical disruption.

“AI has introduced new forms of dependency that evolve faster than traditional governance, procurement, or technology cycles were designed to handle,” said Ana Paula Assis, IBM Senior Vice President and Chair, EMEA and APAC, in the study foreword. “Any loss of control can translate directly into margin pressure, compliance exposure, or outright business disruption.”

Organisations that design AI systems to adapt data, models and infrastructure as conditions change are outperforming peers. The analysis found those with the most advanced AI control capabilities see less AI downtime and protect 55 per cent more operating profit from AI-driven disruptions. Only 7 per cent of organisations surveyed operate at that level.

Cost tolerance for flexibility is notable. Seventy-two per cent of executives say they would accept a 20 per cent cost increase to maintain AI vendors if it improved strategic flexibility. While 73 per cent describe their AI environments as intentionally multi-vendor, the study suggests diversity is driven less by deliberate strategy than by independent business unit decisions (69 per cent), geographic necessity (69 per cent) and legacy complexity (57 per cent).

The full study is available at <https://ibm.biz/ai-sovereignty>

RICOH Scanning Solutions



Fi-8040 – Entry Level 40 Page a Minute A4 Desktop Scanner with LAN and USB Connection



Fi-8150/8170 – Compact, Reliable 50 or 70 PPM A4 Desktop Scanners, Paper Protection, Optimised Image Quality



ScanSnap SV600 – Overhead Style Contactless Scanner, can easily scan business cards, newspapers and magazines up to 30mm thick, scan multiple documents in 1 pass



Fi-7300NX – Secure Wi-Fi Connected Stand Alone 60 PPM A4 Network Scanner



Fi-7600 – Heavy Duty A3 Professional Scanner, 100 PPM, Straight Paper Path, Large Feed Tray, LCD Panel for Easy Operation



Fi-7700 – Similar to the 7600 with A3 Flatbed under the Sheet fed scanner, Mixed document sizes and fragile paper handling on the flatbed in the same batch



Fi-8820 – A3 120 PPM Production Scanner with Automatic Separation Control, Large Touch Screen and both Lan and USB Connectivity



Fi-8930 – Similar to the 8820, A3 130 PPM Production Scanner, Staple Detection, Automatic Skew Correction



Fi-8950 – Top of the Range A3 150 PPM Production Scanner, similar to 8930 but faster and built to scan the largest of volumes every day

RICOH

DOCUVAN
IMAGE and DATA SOLUTIONS

As a SELECT SCANNING PARTNER with Ricoh in Australia, DocuVAN provide access to industry-leading scanning technology backed by our 20+ years of expertise.
Contact info@docuvan.com.au or call on 1300 855 839

The UK says a PDF is not an invoice. France just made it one

By Benson Hendall

Picture two finance managers in the autumn of 2027, one in Manchester and one in Lyon, doing exactly the same thing: exporting an invoice as a tidy, letterheaded PDF and emailing it to a domestic business customer. In Manchester, this is still a perfectly ordinary invoice and will remain one until April 2029. In Lyon, it stopped being one a year earlier.

Since 1 September 2026, invoices between French businesses must travel as electronic invoices through state-registered platforms, and a plain PDF attached to an email no longer qualifies.

Nothing about the file differs. What differs is the definition. Two neighbouring jurisdictions, wrestling with the same VAT-gap arithmetic and the same automation argument, have arrived at incompatible answers to a deceptively simple question: is a document that a human can read, but a machine cannot reliably parse, an invoice? The UK has answered no. France has answered that it can be both, provided the format does both jobs at once.

The UK’s direction was set at Autumn Budget 2025: from April 2029, all VAT invoices for business-to-business and business-to-government transactions must be issued and received electronically. The [consultation response](#) published alongside the Budget is unambiguous about scope: PDFs, Word files, HTML invoices and OCR-scanned images sit expressly outside the definition of e-invoicing. The mandate targets structured, machine-readable invoice data flowing directly between financial systems.

On 23 June 2026, HMRC added the missing technical piece. In its [“Tax update 2026: simplification, modernisation and fairness” policy paper](#), the government confirmed Peppol as the core interoperability network for UK e-invoicing: a decentralised, four-corner model in which suppliers and buyers exchange structured documents through accredited access points, with no central clearance platform and, initially at least, no realtime reporting to HMRC.

The full technical standards are still being co-designed with industry, and a detailed implementation roadmap is due at Budget 2026 in November. But the philosophical commitment is already legible: in the UK’s future regime, the invoice is the data. Any human-readable view of it is a presentation-layer courtesy, generated by software on demand.

What France decided instead

France gets there three years earlier, and by a different route. From [1 September 2026](#), every business subject to French VAT must be able to receive electronic invoices; large and mid-sized companies must issue them from that date, with smaller businesses following on 1 September 2027. Exchange runs through state-registered platforms (plateformes agréées), and the reform admits three core formats. Two of them, UBL and CII, are purely structured XML. The third is Factur-X.

Factur-X is the interesting one. It is a hybrid: a fully

Billing Invoice

COMPANY NAME

Address -

Telephone -

Fax -

Email -

CUSTOMER/CLIENT

Address -

Telephone -

Fax -

Email -

Unit #	Product ID	Description	Per Unit Price	Total Price

Terms & Conditions:

conformant PDF/A-3 document, readable by any human with any PDF viewer, carrying an embedded XML file that machines process directly. The trick is only possible because PDF/A-3 did something its stricter predecessors forbade: it allowed arbitrary files to be embedded within an archival PDF, on the explicit understanding that the PDF remains the authoritative, human-readable rendition of the whole. A Factur-X invoice is therefore one artifact with two native readers. The accountant sees a familiar invoice; the accounts-payable system never touches the pixels and reads the XML. Neither view is a derivative of the other: both are the invoice.

France did not mandate the hybrid; pure UBL or CII are equally valid within its core set. But by placing Factur-X in that set, France made human readability a first-class property of a legal invoice. The UK’s definition excludes it by construction.

Two philosophies of the same document

Strip away the acronyms and the disagreement is philosophical. The UK’s position: an invoice is a record of structured facts, and rendering those facts for human eyes is the software’s job: accurate, current, but existing only when summoned. The French hybrid position: an invoice is also evidence, and evidence should be inspectable by an unassisted human being, today and in 20 years, without depending on the software that produced it.

Readers may recognise the shape of this argument. It is the same gap that separates [a PDF that renders beautifully from a PDF that machines can actually read](#): visual fidelity and machine readability are different properties, and assuming one implies the other is where documents quietly fail.

E-invoicing regulation is that same gap, promoted from an engineering nuisance to a matter of law. The UK resolves it by abolishing the visual artifact from the definition. Factur-X resolves it by binding the two properties into a single file and refusing to choose.

Both positions are coherent. Pure structured data is lighter, less ambiguous, and cannot drift out of sync with a visual layer; there is no visual layer to drift from. The hybrid costs more bytes and more standardisation effort, but it degrades gracefully: whatever happens to platforms, protocols and vendors, the document still opens.

The archive question nobody is asking

That graceful degradation matters most at the point almost no one in this debate is discussing: year six. UK record-keeping rules require invoices to be retained for six years in a form that guarantees access, readability and retrieval. But the readability of a pure UBL or CII invoice is entirely mediated: it means a stylesheet, a renderer, a platform, or an archived viewer that still runs. None of these are properties of the file. They are properties of an ecosystem around the file, and ecosystems have shorter lifespans than tax retention schedules.

A Factur-X invoice, by contrast, is self-archiving by construction. PDF/A-3 exists precisely to guarantee that a document remains self-contained and renderable decades on; that is the format’s founding contract, and the structured payload rides inside the preservation wrapper. It is a curious irony of the UK’s machine-first definition that it makes the invoice easier to exchange in 2029 and potentially harder for anyone to read in 2039. The roadmap due at Budget 2026 would do well to say as much about retention as it does about transmission.

If you invoice across the Channel

For businesses caught between the two regimes, three practical bearings follow. First, know what your jurisdiction calls an invoice: the same file can be compliant on one side of the Channel and legally invisible on the other, and that divergence is live from this September, not hypothetical.

Second, let your counterparties set your real deadline: buyers in mandated markets increasingly ask suppliers for structured invoices as a condition of doing business, whatever domestic law requires and whatever April 2029 permits.

Third, do not discard the PDF; requalify it. In one regime it remains the invoice’s legal body; in the other it becomes the human-facing view of a data record. A demotion in law, but not in usefulness.

The mistake, in other words, is not choosing the wrong format. It is assuming that the phrase “PDF invoice” still means anything at all without saying which side of the Channel you are standing on.

(Both Australia and New Zealand support Peppol e-invoicing, but neither has gone as far as the UK’s full B2B mandate. Both are using government procurement as the lever: B2G capability mandates now (NZ Jan 2026, AU federal targets through 2026), with supplier obligations following via procurement rules rather than tax law.)

Benson Hendall is the Marketing Manager at PDFSmart, a cloud platform designed to simplify everyday document workflows. Originally published [here](#).

UK Retreats on Digital ID Card

The UK government has abandoned its controversial national digital ID card, with reports that incoming prime minister Andy Burnham is dropping the scheme announced by the previous administration. Savings will be redirected to a tax cut aimed at lowering household electricity bills.

The scheme was proposed by former prime minister Keir Starmer to help combat illegal employment and modernise access to services. TechCrunch reported it was projected to cost £1.8 billion over three years. A parliamentary petition against it drew nearly three million signatures, with critics warning a mandatory card could enable surveillance. The plan echoed a national identity card programme abolished in 2011.

The critical distinction is design. The UK proposed a mandatory, centralised identity card, not featured in Australian or New Zealand ID schemes, which do not involve a national ID card or a central database.

With a user’s consent, Australia’s [system](#) verifies existing identity documents online against the official record already held by government agencies. No new identity is created, and no additional personal information is collected, according to the Department of Finance. The design aims to reduce the repeated sharing of ID documents seen in traditional 100-point checks, and with it the fallout from breaches such as Optus, Medibank and Latitude.

Australia has been here before. A proposed national identity card, the Australia Card, was defeated in 1987 after mass public protest, including a rally of more than 40,000 people in Perth and opinion polls that swung from majority support to majority opposition. The current [Digital ID System](#), underpinned by the Digital ID Act 2024, was deliberately built as decentralised verification rather than a single card or register. The contrast matters for the surveillance concerns that sank the UK card. A decentralised model that checks documents against records agencies already hold, without a new database, avoids the single point of aggregation privacy advocates fear.

The federal government committed A\$654.3 million in its 2026-27 budget to scale the Australian Government Digital ID System and the myID platform across more than 255 online services. From 1 December 2026, accredited private providers will be able to join, with charging for business relying parties from 1 January 2027. Individuals will continue to use myID for free.

New Zealand is taking a comparable, framework-based route. Its Digital Identity Services Trust Framework accredits voluntary identity providers rather than issuing a state ID card. A new AML/ CFT [Identity Verification Code](#), which commenced on 1 July 2026, recognises accredited providers as a standalone verification pathway, and NEC New Zealand became the first accredited provider.

Gartner Flags Risks in AI Agents

By 2028, 70 per cent of US healthcare payers will address workflow, data and service automation using no-code agent builders, cutting manual workloads by 50 per cent, Gartner predicts. The forecast signals a broader shift: AI agent development moving from centralised IT projects to business-led automation, with governance the deciding factor between value and risk.

The prediction appears in the Gartner report [Innovation Insight: No-Code Agent Builders Improve Efficiency for U.S. Healthcare Payers](#).

No-code agent builders (NCABs) offer an integrated design and runtime environment to build, publish and manage AI-powered agents without exposing any code to the builder. Gartner says the tools let business teams create agents through visual or conversational interfaces with minimal IT support, providing a sanctioned alternative to shadow AI deployments.

The warnings are as prominent as the promise. Gartner finds many IT and business teams are not fully prepared to implement the tools, facing difficulty defining use cases, insufficiently trained personnel and inadequate governance structures. Deploying NCABs without proper controls can disrupt essential systems, introduce security vulnerabilities and cause service and operational issues. Rapid growth in the agentic AI market - fuelled by vendor hype, inconsistent terminology and what Gartner calls agent-washing - is also making it hard for leaders to understand available solutions.

Gartner recommends a readiness evaluation before implementation, dual-layer governance separating platform controls from agent-specific functions, a well-defined inventory of use cases, enforced guardrails for citizen developers and a phased rollout starting with small, targeted projects. Use cases cited include prior authorisations, member enrolment, provider directory maintenance, appeals and grievances, eligibility verification, regulatory reporting and process documentation.

While the report focuses on US healthcare payers, its governance framework - separating platform-level controls from agent-level oversight, with clear responsibilities for development, security and maintenance - applies to any organisation opening agent-building to business users.

From Task Bots to Goal-Driven Agents

The core leadership question for enterprise AI is no longer “can we build an agent?” but “can we delegate decisions safely at scale?”, according to a new Capgemini point of view on agentic process automation. The consultancy argues organisations that succeed treat agentic automation as an operating-model change rather than a one-time tooling upgrade.

The report, [Agentic Process Automation for Customer First Organizations](#), positions agentic process automation (APA) as the next evolution beyond workflow standardisation and robotic process



automation. Goal-driven AI agents sense, decide and act across customer journeys and enterprise processes, combining AI reasoning with enterprise orchestration and governance.

Governance by design is a central theme. Capgemini argues agents require explicit guardrails - policies, thresholds, approvals and audit trails - with autonomy boundaries defining where agents can safely act and where human judgement must remain. The design goal is what the report calls safe autonomy: automation that is explainable, recoverable and governed through human oversight.

The report notes that after early experimentation, clients are increasingly focused on agentic control planes: governance structures and technical layers that oversee customer flows across front, middle and back office, ensuring auditability, recoverability and accountable decision-making.

Capgemini’s proposed enterprise architecture includes a semantic layer that translates enterprise context for agentic interpretation, an agent fabric connecting tools and systems through managed APIs and MCP, and an agent control plane covering identity and trust, policy governance, safety monitoring, and audit and compliance evidence.

Defined business processes and workflows remain essential, particularly in regulated environments, the report stresses. What changes is the execution model: agents contribute to outcomes alongside humans and deterministic automation, coordinating decisions and actions within defined boundaries while human specialists retain authority for high-risk decisions and exceptions.

The POV sets out industry applications across banking, insurance, telecommunications and the public sector, and a phased delivery framework running from initial scoping through to scaled operation. It concludes that success lies not in deploying isolated AI capabilities but in transforming the operating model to enable orchestrated, customer-centric outcomes.

Firms Blind to Unsanctioned AI

The share of organisations unable to tell whether employees are using unsanctioned AI tools has nearly tripled in a year, rising from 6.3 per cent to 17.6 per cent. For AI agents the blind spot is larger still, with 21.1 per cent of organisations unable to account for unsanctioned agent activity.

The findings come from AvePoint’s third annual [State of AI Report: Scaling Trust, Control, and Readiness in the Agentic Era](#). The study, conducted with Osterman Research, surveyed 750 enterprise leaders with direct responsibility for information management, data security or AI programs across the Americas, EMEA and APAC.

The report found a stark gap between confidence and outcomes. While 82.7 per cent of respondents said they were ‘very’ or ‘extremely’ confident in their ability to prevent unauthorised AI-related data access, 72 per cent of the ‘very confident’ group experienced an unauthorised access incident in the past 12 months. Among the ‘extremely confident’, the figure was 62 per cent.

Security incidents are widespread. In 2026, 89.5 per cent of organisations reported at least one generative AI-related security breach, up from 75.1 per cent the previous year. For AI agents, 88.4 per cent experienced at least one breach in the past 12 months.

The governance gaps are delaying deployments. Nearly nine in 10 organisations (86.9 per cent) delayed generative AI deployments by an average of almost six months, citing data security and governance concerns. The figure for AI agents was almost identical at 86 per cent.

“Nearly half of global employees are already relying on AI

agents weekly or daily, and organizations are deploying agents faster than they are building the foundations required to trust them,” said Dr Tianyi Jiang, CEO and Co-Founder of AvePoint.

“The constraint on enterprise AI is no longer model capability, but whether organizations have built a trust layer: the data visibility, governance, and enforceable control required to scale AI with confidence. Without it, speed of deployment becomes speed of exposure.”

AI is also reshaping the data that governance must cover. On average, 35.5 per cent of enterprise data is now generated by AI assistants, a figure expected to reach 42.1 per cent within 12 months. Meanwhile, 84.1 per cent of organisations manage at least one petabyte of data, and 78.1 per cent report at least half their data is more than five years old.

Adoption continues to accelerate. Some 46.9 per cent of employees rely on AI agents daily or weekly, and organisations anticipate agents will replace more than a quarter of human work within 12 months. Reducing headcount ranked last among reasons for adopting agents, with return on investment measured instead through reduced manual effort and faster processes.

Investment is following the gaps. Securing data used for AI training was the top-rated future priority at 79.5 per cent. Third-party governance tools that monitor agent actions topped the planned investment list for the next 12 months, capabilities at the core of the emerging AI Agent Management Platform category defined by Gartner.

“Trust in AI cannot be measured by confidence alone,” said John Peluso, Chief Technology Officer at AvePoint. “It requires operational foundations: visibility into what AI systems are doing, enforceable governance over the data they consume and create, and the ability to audit and correct outcomes when something goes wrong.”

Insurer RACQ Bets On Adobe Agentic AI

RACQ has signed a five-year partnership with Adobe and Deloitte Digital to transform how it serves more than 1.7 million members. The deal uses a flexible “Lighthouse” model that gives RACQ early access to Adobe’s latest AI tools. Adobe says it is the first such arrangement in Asia Pacific and Japan.

The Queensland mutual supports members across motoring, insurance, banking, batteries, solar and travel. The agreement aims to improve personalisation, content production and automation across those services. RACQ will expand its use of Adobe’s Customer Experience Orchestration and Creativity and Productivity portfolios. A governance and learning framework will guide adoption and target high-value use cases.

While both companies framed the deal around marketing and member engagement, several named products point to broader operational use inside an insurer.

Adobe Experience Manager and Acrobat Studio handle document-intensive work. Insurers typically use such tools for policy schedules, product disclosure statements, renewal notices and claims correspondence.


Experience Manager Forms commonly powers digital claims lodgement and online policy applications. That positions the technology well beyond campaign content. Adobe Realtime CDP Collaboration unifies member data across RACQ’s six product lines. A single member view supports service, retention and risk teams, not only marketers.

Adobe Experience Platform Agent Orchestrator coordinates AI agents. In insurance settings, that can mean triaging claims queries, routing member requests and drafting standard correspondence.

Adobe Workfront manages work and approvals. Regulated insurers often use it to track compliance reviews and sign-offs on member-facing documents.

Adobe Firefly Creative Production and Creative Cloud for Enterprise round out the content tools. Adobe Ultimate Success with Integrated Services provides governance and technical support.

David Harrys, Chief Member and Growth Officer at RACQ, said the partnership equips RACQ to help members live and move safely, securely and sustainably as mobility, energy, home and climate needs converge. Deloitte Digital will lead the end-to-end implementation of Adobe’s solutions and support RACQ’s broader digital transformation.



Building an enterprise AI governance framework in seven phases

By Konstantinos Bardavouras

Somewhere in your organisation there is a beautifully formatted AI policy that nobody has opened since the day it was approved. I've read a few of them. They tend to be thorough, well-intentioned and completely disconnected from the systems they claim to govern. Governance that lives in a document is decoration. What works is treating it as a **system you build, document and hand off - in that order**. This is the model I use, compressed to seven phases.

01 Mandate

Before any register, control or committee, someone senior must own the outcome. With their name, not their department's. Governance without decision rights is advice, and advice loses every argument with a delivery deadline. I've watched strong frameworks evaporate the first time a flagship project asked for an exception and nobody in the room had the authority to say no.

Phase one is getting three things in writing: an accountable executive sponsor, a definition of what counts as an AI system in your organisation, and who can pause one. That second item sounds bureaucratic until you try to answer it. Is the copilot inside your CRM an AI system?

The scoring model a vendor runs on your data? The agent someone wired up last month to triage tickets? If your definition doesn't cover models you build, agents you assemble *and* AI embedded in the software you buy, it will be obsolete before the ink dries. And if you cannot get the mandate at all - stop. Everything after it is theatre.

02 Inventory

You cannot govern what you have not found, and there is always more than anyone admits. Build a register of every AI use case **and every agent**: owner, purpose, the data it touches, its autonomy level, the systems it can act on.

The uncomfortable part is shadow AI - the unofficial copilots and connectors nobody declared. Expect to find them; every tenant has some. Treat what you find as a demand signal rather than a crime scene. People built those things because something in their job was slow or painful, and the register is how that need gets a sanctioned path instead of a hidden one. Punish the first discovery and you guarantee you'll never make another.

03 Classify

Risk-tier every entry: where it lands under the EU AI Act, what breaks when it fails, who gets hurt, whether the harm can be undone, how sensitive the data is. My favourite classification question is blunt: **what happens if this is wrong and nobody notices for a week?**

The honest answers range from "nothing" to "we make the news", and that range is exactly what the tiers must capture.

Keep the tiers few - three or four, not nine - and predefine the obligations for each. Classification is what turns governance from an argument about opinions into a lookup. The tier tells you what is required, nobody negotiates controls case by case, and the review meeting takes twenty minutes instead of a quarter.

04 Control

Controls should be proportional to tier and enforced where the traffic actually flows. Not in a document, not in a training slide. In practice that means one controlled front door for AI workloads: a gateway handling routing, quotas and content safety at the edge, identity on every call, hard data boundaries, and policy expressed as code wherever the platform allows it.

A control that depends on people remembering it is a suggestion. A control enforced by the platform is a fact.

You want as many facts and as few suggestions as your architecture can afford, because eighteen months from now nobody remembers the suggestions.

05 Evidence

At audit time, a control without evidence does not exist. It doesn't matter that the guardrail fired; if nothing recorded it firing, you cannot prove it was ever there.

So every tier's obligations must produce artifacts continuously and automatically: evaluation pipelines on every release, red-team results, groundedness checks for anything that answers from enterprise content, logs that show the boundaries holding. The test is simple and slightly cruel: if assembling evidence for an audit takes a quarter of manual archaeology, phase five was skipped. No matter what the policy says.

06 Document

Now write it down. Not the compliance novel - the operating model people will actually use: short policies, runbooks, a decision log, the risk register itself. My test for "done": a new engineer can find out what they are allowed to build, and how, without booking a meeting with the person who designed the framework.

This is also where aligning to ISO/IEC 42001 and NIST AI RMF earns its keep. It costs little while you are structuring the material anyway, and it turns every future audit from an excavation into a mapping exercise. Do it now, or pay for it annually.

07 Hand off

Here is the part I have to argue for most often: the architect who built the framework is the wrong permanent owner for it.

Partly it is a conflict of interest - the person who built the system should not be the only one who can say no to it. Mostly it is a single point of failure: governance that lives in one person's head leaves in their notice period. And there is a quieter failure mode too, the flattering one. It is pleasant being the person everyone has to ask. Resist it.

Accountability transfers to named, permanent owners - risk, security, data, product - each holding the piece that matches their mandate. The architect stays close, as an advisor for the genuinely hard calls, and the framework keeps running on the days the architect is on a beach. **Governance is an operating function, not a project.** Projects end.

None of this makes AI governance easy. It makes it operable, which is the achievable kind of easy. The standards - ISO/IEC 42001, NIST AI RMF, the EU AI Act - give you the skeleton and the vocabulary; the seven phases give you the order of operations. Build, document, hand off. Then let go properly, on purpose, with your name off the approval chain. That is the only version of "done" that survives you.

Konstantinos Bardavouras is an AI Solution Architect. Article originally published [here](#).

Ancestry.com Bets \$US450M on Digitisation

Ancestry.com, the world’s largest online family history platform, has cut a record digitisation process that once took nine months of manual indexing down to under nine days, using proprietary handwriting-recognition AI to help push its collection past 71 billion records and restore subscription revenue growth after years of stagnation.

The US-based company, which operates in more than 30 countries and gives subscribers access to census, birth, marriage, death, immigration and military records alongside DNA testing and family tree building tools, has become one of the largest repositories of digitised historical records in the world.

It was acquired by private equity firm Blackstone for US\$4.7 billion in 2020.

Ancestry CTO Sriram Thiagarajan gave [Authority Magazine](#) a detailed account of the pipeline in June. Before 2021, records were scanned and then manually indexed through outsourced operations, a process taking almost nine months at roughly ten times the cost.

Proprietary AI handwriting recognition compressed that to under nine days, and Ancestry has more than tripled its rate of content growth year over year.

“At Ancestry, we see AI as an amplifier of human capability, not a replacement,” Thiagarajan said.

“In our industry where we lift insights from historical records and validate deeply personal family histories,

real human involvement is critical.”

He said human review remains embedded in the workflow to guard against AI hallucination and to preserve historical accuracy and cultural sensitivity.

The turnaround has also been detailed by Ancestry’s CEO. Howard Hochhauser told [Semafor](#)’s CEO Signal in April that the company plans to invest US\$450 million over the next 10 to 15 years in digitising new records and applying AI to extract more accurate information from previously scanned material.

Hochhauser, a former CFO who became CEO in early 2025, told Semafor the company had drifted from its core customer under a strategy of making the site more social.

“Revenues flattened out, customer satisfaction declined,” he said. Refocusing on content and storytelling has since delivered five consecutive quarters of subscription revenue growth and 13 months of rising customer satisfaction scores.

The content engine is archival. Hochhauser described photographing paper voting registers in the basement of Manhattan’s City Hall, after which “technology takes over and AI will take that handwritten script and make it searchable for everybody. Every day, we’re writing 10 million new records to the site.”

The company is investing in birth, marriage and death records in the US, UK, Canada and Australia, with Italy and Poland next.

The generative layer sits on top of the structured extraction. An AI Stories feature, launched in December, turns static documents such as census records into narrative and audio stories; more than 500,000 unique users have engaged with it and almost a million audio stories have been played.

AI translation has also enabled rapid international expansion, with the site, content and marketing campaigns launched in Italy within a month.

Technology publication [Let’s Data Science](#), analysing the interviews, called Ancestry a case study in how domain-specific OCR, generative storytelling and human-in-the-loop validation can compound into a defensible data advantage over a decade.

It flagged error rates on non-English historical scripts, dataset provenance and cross-border privacy obligations as the issues to watch as the approach scales internationally.

Hochhauser also noted the limits customers place on automation. “It’s somebody’s hobby, the joy is in the doing and discovery,” he told Semafor. “You don’t want your AI assistant to do your crossword puzzle for you; that’d be wholly uninspiring.”



Kapish

Empowering Secure Technology Solutions



Talk to us today to find out how our suite of products and services can help you get the most out of Content Manager.



Call 1300 KAPISH | info@kapish.com.au | kapish.com.au



Spy agencies say AI can help combat AI cyber risks. But don't forget the basics

By Toby Murray, The University of Melbourne
Cybersecurity agencies of Australia, Canada, New Zealand, the United Kingdom and the United States have issued a [call to action](#) for cyber defenders. The message was clear: artificial intelligence (AI) is a powerful weapon for cyber attackers; defenders must act urgently to improve their cyber defences.

There is much hype and uncertainty surrounding AI and cybersecurity right now. This latest statement came little over a week since the US government caused frontier AI provider Anthropic to [block access](#) to Mythos and Fable, its most advanced AI technology, over fears they might be misused by foreign adversaries to [attack US government systems](#).

In this torrid environment, it's important for cyber defenders to look past the noise and prioritise what is truly important in protecting their systems.

The joint statement was issued by the heads of the national cybersecurity agencies of the [Five Eyes](#). It warns that AI is dramatically shifting cyber risk and spells out how defenders must act to secure their organisations.

It notes how powerful AI is already helping adversaries carry out more sophisticated attacks more quickly.

One way this is happening is through automated vulnerability discovery and exploitation. No software is perfect. Adversaries leverage subtle design or implementation flaws in a system's software to break into that system. They then take control of it and use it as a staging ground to launch further attacks.

This is why it's so important for cyber defenders to keep up to date with deploying software patches. These are small modifications to system software that close off known vulnerabilities.

AI is enabling adversaries to find flaws orders of magnitude faster, as well as to work out how to exploit those flaws to [carry out attacks](#).

For this reason, the Five Eyes statement warns that AI is dramatically shrinking the time between when a vulnerability is first discovered and when it is first exploited in an attack. Defenders can [no longer afford](#) to wait weeks before deploying software patches.

What can defenders do?

The Five Eyes report notes cyber fundamentals are crucial and encourages organisations to use AI to boost defences. But deploying AI without first investing in cybersecurity basics would be a mistake.

The cyber defenders who will be able to weather the AI storm will be those who already have mature practices. They know exactly what assets they need to protect, which systems in their organisation are exposed to attack, and

what defences are in place to protect exposed systems. They also know to measure defence effectiveness and determine where defences are missing.

They also use evidence-based processes for tracking known vulnerabilities in their systems and prioritising which are most important to patch. These are backed up by reliable processes for rapidly testing and rolling out software patches, as well as for responding to cyber breaches and incidents.

When AI makes finding software vulnerabilities cheap, the next generation of software needs to be engineered to be [secure by construction](#).

Working out the [best methods](#) to do this is what I have devoted my research career to.

Before reaching for AI, defenders should first invest in their fundamentals. Otherwise, they are effectively deploying a robot guard dog to defend an unlocked door.

The role for AI in cyber defence

This doesn't mean AI can't play an important role for cyber defence – just that it should augment rather than replace strong cyber fundamentals.

AI benefits attackers and defenders alike. An AI model that can help attackers find software vulnerabilities can also help defenders fix those same vulnerabilities.

AI that can automatically exploit software vulnerabilities is just as useful to defenders in helping them to confirm their software has been correctly patched. AI that can map and discover sensitive assets within a computer network is useful for both offensive and defensive purposes.

This is why it's so important that defenders have access to AI capabilities, so they can be leveraged to harden and protect systems before that same AI is used to attack them.

Can regulation help?

Working out how to balance the competing benefits and risks of new cybersecurity technology is nothing new.

In the 1990s, society [grappled with](#) how to regulate the encryption that protects online communication from adversaries but also allows them to avoid law enforcement.

In the 2000s the rise of cyber exploit kits allowed defenders to better test their systems but also enabled any disaffected teenager with an internet connection to become a "script kiddie" hacker, leading to [arms controls debates](#) a decade later.

The 2010s gave us blockchain technologies such as Bitcoin and other cryptocurrencies, which were built on defensive cyber technologies but whose lasting legacy remains the rise of ransomware attacks and online illicit marketplaces.

The rise of AI presents a similar dilemma for regulators.

A blanket export ban on advanced AI models is [likely to be counterproductive](#). Open-source AI models such as [DeepSeek](#) lag only months behind the most advanced models of OpenAI and Anthropic. Recent research suggests that much of that gap [can be closed](#) by pairing less powerful AI models with complementary technologies.

Defenders should therefore assume their adversaries already have access to AI on par with that used for cyber defence. Only by investing in strong foundations can they hope to escape the cat-and-mouse AI cyber arms race.

[Toby Murray](#) is Professor of Cybersecurity, School of Computing and Information Systems, [The University of Melbourne](#). This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Beyond compliance: what AML reforms reveal about information governance

By Garry Valenzisi

From 1 July 2026, Australia's expanded Anti-Money Laundering and Counter-Terrorism Financing regime will extend to a broader range of professions, including lawyers, accountants, real estate agents and other designated non-financial businesses and professions.

For many organisations, the immediate response will be to focus on the new rules: what needs to change, who is responsible and what evidence must be produced. That is understandable. But the reforms also raise a broader and very practical business question: if you were asked to prove compliance tomorrow, could you quickly find and trust the information you need?

This is where AML reform becomes more than a regulatory issue. It becomes an operational one.

Globally, businesses are already feeling the pressure. [More than three-quarters of organisations say compliance complexity is negatively affecting their business](#), with rising costs, increasing resource demands and greater operational risk.

For business leaders already managing tight margins, lean teams and competing priorities, another layer of compliance can feel like yet another burden.

But it can also be a useful trigger to fix a problem many organisations already know they have: information that is hard to find, inconsistently managed or spread across too many systems.

The real compliance challenge is control

[Iron Mountain research shows](#) a strong data foundation is a top strategic priority for many Australian businesses.

Yet globally, organisations continue to struggle with the basics of information management. [Almost six in ten executives cite data availability and quality as a major barrier to reporting](#), while only one in five say they have fully validated their data.

For an executive team, this is not just a technology issue. It affects how quickly a business can respond to a regulator, satisfy an auditor, answer a customer query, manage risk or make a confident decision.

In many organisations, information sits across physical records, legacy systems, shared drives, email inboxes, cloud platforms and business applications.

Different teams may have different ways of storing, naming and retaining records. Ownership can be unclear. Duplication is common. Retention rules may be applied inconsistently.

Most of the time, these issues sit quietly in the background. But when scrutiny arrives, they quickly become visible.

Experience in markets such as the UK and New Zealand shows the risk of treating compliance as an administrative exercise.

Organisations have not only had to understand new AML requirements; they have also had to demonstrate that records are complete, decisions are documented

and risk assessments are supported by reliable data.

For Australian businesses, the lesson is simple: compliance depends on information you can trust.

Why audit readiness is a business priority

Being audit-ready is not about preparing for one review or responding to one regulatory request. It is about having a consistent, defensible way to manage business information every day.

That starts with knowing what information exists, where it is stored, who owns it, who can access it, how long it should be kept and when it should be securely destroyed. These may sound like back-office details, but they have very real business consequences.

When information is well managed, teams can respond faster. Compliance costs are easier to contain. Risk is easier to monitor. Decisions are based on better evidence. Customers and partners can have greater confidence in how information is handled.

When information is poorly managed, the opposite is true. Poor records management has been linked to [60 per cent of data breaches](#), highlighting the connection between governance and business risk.

Fragmentation remains one of the biggest barriers. Disconnected systems and inconsistent processes create blind spots that make it harder to enforce policy or show accountability.

As regulatory expectations rise, [only 43 per cent of organisations have formal governance frameworks in place, while 80 per cent report inconsistent practices across environments](#).

The issue is rarely that businesses do not have enough data. Most have more than they can comfortably manage. The issue is whether they have control over it.

Turning reform into a practical advantage

For organisations preparing for expanded AML/CTF obligations, the message is clear: strong information governance is essential to effective compliance.

But this should not sit with one person or one team. Legal, compliance, finance, operations, IT, records management and executive leadership all have a role to play. Policies matter, but so do everyday behaviours: how teams save documents, classify records, manage access, dispose of outdated information and respond when evidence is needed.


The expansion of AML/CTF regulation is a pivotal moment for Australian businesses. Regulators will not only be interested in whether policies exist. They will want to know whether organisations can demonstrate control through accurate, accessible and well-governed information.

For business leaders, that makes information governance more than a compliance task. Done well, it can reduce risk, improve efficiency and create greater confidence in the way the organisation operates.

Safe, compliant business starts with information you can trust.

Garry Valenzisi is Vice President & General Manager at Iron Mountain ANZ.

Smart Scanning Solutions for Any Document Type



Book Scanners



Flatbed Scanners



A3 Production Ricoh



Wide Format Scanners



XINO S700 Series

DocuVan is a Distributor and Reseller of higher end scanning equipment.
We can supply, install, train and support you in operating 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.



BookTEK®

WideTEK®

RICOH
PLATINUM BUSINESS PARTNER

Janich & Klass
Computertechnik

DOCUVAN
IMAGE and DATA SOLUTIONS

MAKE
ENQUIRY

Email info@docuvan.com.au or call on 1300 855 839



Active Metadata: The Key to Content Observability

By Ole Olesen-Bagneux

Employees and customers increasingly encounter AI-generated content via emails, chatbot responses, learning materials, and documents. But enterprises have no way to know how that content, also known as unstructured data, is performing. To improve content over time, it makes sense to observe the human reactions to it. Is it clear? Is it trustworthy? Is it driving the desired actions?

AI feeds on large volumes of unstructured data, so to be successful, enterprises need to manage it at a level of detail and precision not seen in the past.

The best approach to content observability leverages active metadata.

What is Active Metadata?

Gartner popularized the notion of active metadata in the early 2020s as a strategic mindset for enterprise metadata management. While descriptions vary slightly,

active metadata essentially has four characteristics:

■ **Always on** means that metadata should be fresh. Although true realtime metadata is rarely worth the costs, metadata should provide relevant, updated context for whatever purpose the metadata serves. Let's take the example of databases, which you don't scan once, but repeatedly. To be reliable, the metadata must be updated.

■ **Intelligent** refers to improving metadata regularly with AI. With AI, generating business glossaries, data descriptions and tags is more efficient than ever. The patterns that emerge as people search and find databases in a data intelligence platform must feed back into the business glossaries to make it easier to find those databases in the future.

■ **Action-oriented** means that change notifications should be pushed to end users to make them engage and take action, instead of waiting for the metadata to be discovered. If you often search for a particular database and a major change occurs in it, you are automatically notified of the change.

■ **Open by default** is the idea that the content of metadata management solutions should point towards other solutions through API's and, more recently, through MCP and A2A. As a result, metadata is surfaced in tools the users already work with, so it is useful to them. When metadata about a given database is not kept in the data intelligence platform, but you can see it elsewhere, such as in a business intelligence solution.

Active metadata was built for structured data, which was the main focus of past enterprise data management approaches. This framework translates surprisingly well to unstructured, AI-generated content.

What is Content Observability?

Content observability is still being defined, but is a modification of data observability. Data observability is the monitoring of structured data as it travels through the organization's pipelines. What is measured is not only the quality of the data when it is stored somewhere, but the state of the data as it moves. Are the pipelines solid, or do they break? Is the data quality upheld from source to destination? Is sensitive data filtered out correctly so it is not visible to all end users of a BI solution?

Content observability takes the principles of data observability and applies them to AI-generated content. Instead of examining how structured data moves, content observability examines unstructured data that moves and is interpreted by humans. Is the content deteriorating as it is extracted from a model and exposed to a user? What

unstructured data sources fed the model? What is the psychological effect of the content - is it precise, confusing, vexing, or boring?

Content observability is in its infancy, but growing in importance as AI-generated content is used to support an increasing amount of business processes. Why? Because our reaction is vital to improve the content your colleagues will be introduced to in the future. Collecting and measuring human reactions - that's content observability.

Intuitively, content observability would look at the unstructured data itself. But what if we moved beyond the data toward metadata? We could improve the content by using metadata strategically. Beyond improving the phrasing of content, it could select data sources more fit for the purpose of content creation.

That's the hidden potential in combining content observability with active metadata.

Combining Active Metadata and Content Observability

Using active metadata to improve content observability forms something enterprises have been missing: a feedback loop that makes AI-generated content smarter over time, grounded in how real humans actually respond to it. What active metadata adds is not only the direct refinement of the AI models, but improves the selection of the data sources that feed those models.

If we apply the four characteristics of active metadata to content observability, we see the benefits.

■ **Always on:** The AI-generated content presented to users would be based on fresh metadata. When guided by better sources and more exact contextual understanding of the human recipient, the content will be better received.

■ **Intelligent:** By creating a feedback loop, the metadata improves as users interact with the AI-generated content. Content observability provides evidence that informs those improvements.

■ **Action-oriented:** Pushing active metadata into the systems where end users already work and into their conversations helps humans engage more effectively with AI-generated content.

■ **Open by default:** Metadata should be accessible and usable without any limitations. Making metadata available through various mechanisms, such as MCP, is a prerequisite to enable content observability.

The use cases for combining active metadata and content observability are many and would altogether improve metadata and human/AI interactions. Consider a financial services firm deploying a customer-facing chatbot. Are customers abandoning conversations mid-way? Asking the same question repeatedly? These are signals that the content isn't landing. Content observability captures those signals and active metadata feeds this back into the data sources shaping the chatbot's responses to improve how answers are phrased and which data sources that inform them. In a regulated industry, that feedback loop is as much a risk management tool as a quality one.

The same principle applies across domain-specific email generation, AI-produced learning materials, and standard operating procedures.

Content observability tells you how your AI content is landing. Active metadata helps you do something about it.

Ole Olesen-Bagneux is VP, Chief Evangelist at [Actian](#). This article originally appeared in [Insight Jam](#), an enterprise IT community that enables human conversation on AI.

Most enterprises lack data foundation for AI

Nearly 80% of enterprise IT leaders say limited data access is constraining their AI and data initiatives, despite most organisations reporting active AI integration across core business processes, according to a new Cloudera survey. The Data Readiness Index, drawn from 1,270 IT leaders across the Americas, Europe, and Asia-Pacific, found a stark gap between perceived preparedness and actual capability.

While 96% of respondents said their organisation has integrated AI into core business processes, and 85% said they have a clear data strategy, the majority still report significant data access limitations.

Cloudera describes this as an 'AI readiness illusion' - a belief that organisations are prepared to scale AI even when critical data challenges remain unresolved.

The top barriers to AI return on investment cited by respondents were data quality (22%), cost overruns (16%), and poor integration into existing workflows (15%). Nearly three-quarters (73%) said performance constraints have hindered operational AI initiatives.

The survey found that only 18% of respondents said their data was fully governed. While 71% said most of their data is governed, Cloudera argues that partial governance is insufficient to support enterprise-wide AI.

The report also reveals significant variation across industries. In telecommunications, 54% of respondents said they have full visibility into where their data resides and 51% said they can access all data at any time. In financial services, those figures drop to 30% and 24% respectively. In the public sector, only 31% have full data visibility and just 16% report unrestricted access.

"Enterprises aren't struggling to adopt AI, they're struggling to operationalize it beyond experiments," said Sergio Gago, Chief Technology Officer at Cloudera.

"AI is only as effective as the data that fuels it. Without seamless access to all their data, organizations limit the accuracy, trust, and business value that AI can deliver."

<https://www.cloudera.com>

Vic Councils' Widening AI Governance Gap

Almost two thirds of Victorian local government officers are using artificial intelligence without clear governance behind them, a new survey has found. The gap exposes councils to privacy, accountability and compliance risk.

The finding comes from [Bridging the Gap: AI Risk and Governance in Local Government](#), a report by law firm Maddocks. It draws on responses from 337 officers across 75 Victorian councils, surveyed in September and October 2025. A follow-up pulse survey of 100 officers in June 2026 confirmed the trends.

The report describes a sector where AI adoption is real but tactical. In the pulse survey, 92 per cent said their council used AI in some form, up from 68 per cent in late 2025. Yet only about a third across both surveys believed their council had a governance framework in place. Most use remains individual rather than institutional. In the pulse survey, 54 per cent said AI was used mainly for individual productivity such as drafting. Just 2 per cent said it was embedded in business

processes or service delivery. Drafting and writing was the top task, mentioned 186 times in open responses.

Privacy and data protection was the biggest risk officers associated with AI. It was cited by 95 per cent of respondents in the initial survey and 62 per cent in the pulse survey. Quality and reliability of outputs narrowly led the pulse survey at 63 per cent, ahead of governance and policy gaps at 49 per cent.

The report notes that officers are proactively self-imposing limits, such as avoiding uploading sensitive data and closely reviewing outputs. It links this caution to a lack of guidance rather than resistance to the technology, concluding that officers "are not anti-AI, but rather anti-risk".

Uncertainty runs through the data. On average, 67 per cent were unsure whether their council had received legal or compliance advice on AI use, council data or output.

EDRM Overhauls Governance Model

The Electronic Discovery Reference Model (EDRM) has released version 4.1 of its Information Governance Reference Model (IGRM), the first major revision since 2012. The update arrives as organisations confront AI adoption, expanding privacy obligations and growing pressure to coordinate information decisions across departments. The IGRM is a reference framework that promotes policy and process integration across the people who create, manage and govern enterprise information.

Its stated purpose is to support conversations about information governance among diverse stakeholders, giving them a common language for decision-making. The model identifies seven stakeholder groups: business, IT, security, records and information management (RIM), legal, risk and privacy.

EDRM describes the IGRM as an intellectual framework rather than an implementation manual. It maps how information moves through a lifecycle, from creation or receipt through use, protection, retention and transfer, to eventual disposal. At its centre sits a perpetual policy and process cycle. The model treats governance as never finished, evolving through the interplay between policy and real-world implementation.

Key enhancements include expanded guidance on unified governance structures and clearer articulation of stakeholder viewpoints and interdependencies, according to EDRM. The revision also offers frameworks for balancing value and risk, and strategies for operationalising governance at enterprise scale.

The model addresses the intersection of information governance and data governance as AI drives more business decisions. The accompanying User Guide notes the two disciplines are distinct but complementary.

Data governance focuses on structured data, quality and metadata within IT. Information governance is broader, covering unstructured and physical information across all seven domains.

The detailed User Guide explaining the model was completed by a reassembled project team in 2025 and has only now been published.

The IGRM v4.1 and its User Guide can be downloaded from the EDRM [website](#).



FREE WEBINAR • ON DEMAND

Overcoming Content Chaos: How AI Transforms Unstructured Data into Actionable Insight

~40 minutes

Access Now →

Your enterprise data has answers. The problem is finding them.

Contracts, invoices, case files, scanned documents, emails — most organisations are sitting on vast reserves of unstructured content that their systems simply can't read, search, or act on. The result? Slower decisions, compliance exposure, and genuine business value buried in digital filing cabinets.

This practical session from Hyland is designed for information and content management professionals who want to understand where AI is delivering real results in document-heavy environments — and how to get there without the disruption of a full platform overhaul.

What you'll take away

- A clear-eyed view of where AI adds genuine value in content and document-heavy workflows
- How intelligent document processing can surface insight from unstructured content at scale — across government, financial services, legal, and healthcare
- Why integration with your existing ECM, ERP, or line-of-business systems matters more than the AI model itself
- Real examples of organisations automating document classification, data extraction, and intelligent routing — using existing infrastructure

View On Demand Now >>



Hyland™



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, Information'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, Information 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 realtime discovery, automated classification, ROT clean-up, compliant retention, and secure management of sensitive information, without compromising compliance. Information 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.information.com.au | info@information.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



EncompaaS is a global software company specialising in information management, powered by next-gen AI. Leading corporations, government departments and statutory authorities trust EncompaaS to govern and optimise information that resides within on-premises and multi-cloud environments. Organisations are empowered to solve information complexity, proactively address compliance and privacy risk, and make better use of data to act strategically at pace. EncompaaS is distinguished in the way the platform utilises AI to build a foundation of unparalleled data quality from structured, unstructured and semi-structured data to de-risk every asset. From this foundation of data quality, EncompaaS harnesses AI upstream to unlock knowledge and business value that resides within information. EncompaaS maintains a robust partner ecosystem, including global consulting and advisory firms, technology partners, and resellers to meet the diverse needs of highly regulated organisations.

encompaas.cloud | enquiries@encompaas.cloud



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

Kapish Partners With Ardoq

Australian information management specialist Kapish has formed a strategic partnership with Ardoq. The deal adds Ardoq's enterprise architecture platform to Kapish's portfolio for customers across Australia and the region. The agreement lets Kapish offer Ardoq's AI-first platform alongside its own delivery and advisory services. Kapish says the combination helps customers cut complexity and align technology spending to business strategy.

Enterprise architecture maps how strategy, applications, data, processes and people connect. Ardoq says its platform replaces static documentation with a living view, using dynamic visualisations and connected data updated in realtime. Kapish says the platform suits work such as digital transformation, application rationalisation, technology lifecycle management, compliance and risk reduction.

"Our partnership with Ardoq represents a deliberate step forward in how we help our clients manage complexity and drive better technology decisions," said Ryan Harris, managing director of Kapish.

"By combining Ardoq's modern, data-driven enterprise architecture platform with Kapish's market-leading expertise and proven delivery capability, we are enabling organisations to move faster, reduce risk, and gain clear line of sight from strategy through to execution."

<https://kapish.com.au>

Automated Redaction Cuts Disclosure Effort

Castlepoint Systems has added automated redaction to its information governance platform. The capability extends support across the disclosure and eDiscovery lifecycle for organisations preparing information for release.

The capability automatically redacts PDF and Microsoft Office documents, replacing sensitive text with black blocks, highlighted spaces or stars. Where human judgement is required, documents can be flagged for manual review before redaction is applied. Castlepoint says the feature helps prepare information for disclosure while protecting content that must remain confidential. It aims to reduce manual effort, improve consistency and lower the risk of human error across the disclosure process.

The Australian vendor frames redaction as central to freedom of information requests, court proceedings, audits, investigations and eDiscovery. In those processes, personal information, commercially sensitive material and security-related content often sit alongside information that should be released. A single missed sentence, it notes, can lead to privacy breaches, regulatory penalties and reputational damage.

Castlepoint argues manual redaction no longer scales, given that a single investigation or FOI request can involve thousands of emails, reports, spreadsheets and attachments. Automated redaction processes sensitive information consistently, while documents needing judgement can still be sent for manual review.

The company positions redaction as part of wider information governance rather than a standalone task. It says its existing metadata, classification and search capabilities already provide the context needed to understand information. Applying those insights to redaction creates a more consistent and defensible approach to disclosure.

The announcement follows Castlepoint's recent IRAP assessment at the PROTECTED level. The company says the assessment demonstrates its commitment to secure and auditable information governance for government and highly regulated industries.

Automated redaction is available immediately as part of the Castlepoint Systems platform.

<https://www.castlepoint.systems>

eGain Grounds AI in Knowledge Hub

Enterprises can now ground Microsoft Copilot, Anthropic Claude, Google Gemini and Cursor in a single governed knowledge source through new connectors from eGain. The connectors link the AI platforms to eGain AI Knowledge Hub, the company's AI-powered knowledge management platform. Every model, agent and agentic developer environment can then draw on accurate, up-to-date knowledge rather than fragmented and outdated sources.

eGain pointed to research from MIT and analyst firms showing that enterprise AI initiatives fail at high rates when they lack a governed knowledge foundation. When any AI system pulls from fragmented, ungoverned sources, the company said, it produces contradictory outputs, creates compliance exposure and stalls AI adoption.

The new connectors are part of the broader eGain AI Knowledge Connector family, which supports the Model Context Protocol, the emerging standard for connecting AI agents to enterprise systems. eGain also offers pre-built connectors for agentic developer environments including Windsurf, VS Code and Kiro, and supports any MCP-compatible platform through its open architecture.

Unlike basic retrieval approaches, eGain said, it delivers certified answers with source citations, giving every connected model, agent and IDE access to verified, traceable knowledge. The same governed knowledge that powers customer service and employee productivity can now inform how enterprise software is built.

The MIT research eGain cited has been widely discussed for its finding that most enterprise AI pilots fail to reach production.

<https://www.egain.com/>



How to Contain Rogue Agents

Delinea has released new runtime authorisation for AI agents, enforcing policy on what an agent does inside a session before it acts, not just how it connects. The identity security vendor said the capability lets organisations evaluate, allow, block or escalate agent actions in realtime. It works across databases, cloud environments, Kubernetes clusters, SSH hosts and Model Context Protocol-connected systems.

The launch targets a gap opening up as automated systems gain access to business-critical resources. AI agents now operate autonomously across production databases, SSH hosts, Kubernetes clusters, cloud consoles and MCP servers. Delinea said most identity security approaches either verify a connection when a session opens, then audit or revoke access after the fact, or can only see traffic routed through their own server.

Neither approach, the company argued, enforces policy on what happens inside the session, action by action, before it executes. A single AI session can carry dozens of tool calls, each with a different risk profile. Agents never hold standing credentials. Delinea injects a credential just-in-time at the

moment of access, scoped to the task rather than inherited from the person who deployed the agent, and revokes it automatically when the task completes. Because the credential is never exposed to the agent, the company said, there is no lingering exposure between sessions for an attacker to exploit.

The platform evaluates and enforces policy on each tool call individually, allowing it, blocking it, or pulling in a human before it executes. When an agent misbehaves, Delinea said, the blast radius stays contained. The platform also identifies whether a connection is agent-driven or human-driven before granting access, applying agent-specific policy so the right controls are in place before the first action runs.

Delinea said it authorises and records each connection through a single control point, across every protocol agents use, giving security teams the same visibility and a complete, defensible audit record regardless of how an agent was deployed.

CEO Art Gilliland said an organisation could have perfect credential hygiene and still have an agent tear through a production environment in milliseconds. Recording what happened or revoking access after the fact does not stop that, he said, but enforcing policy on the action as it runs does.

<https://delinea.com/>

Elastic AI Models Fully On-Premise

Elastic has announced a new offering that provides enterprise-grade data extraction and semantic search with no internet connection or third-party AI services required. It is aimed at regulated industries, air-gapped environments and teams wanting full control over data, cost and performance.

Organisations in regulated or disconnected environments can now run Jina AI models entirely on their own infrastructure through Jina On-Prem, released by Elastic.

Jina AI is a Berlin-based company founded in 2020 and known for its open-source, search-grade embedding and reranking models, used widely for neural search and retrieval-augmented generation. Elastic acquired Jina AI in October 2025, and Jina On-Prem is the first major packaging of those models for the Elastic stack.

The underlying models power the machinery beneath modern search. They turn documents into embeddings, which are numerical representations of meaning, then rerank results for relevance and read and extract text from files. Semantic search built on embeddings retrieves records by concept rather than exact keywords, so a query can surface a relevant document even when it shares no words with the search terms.

Jina On-Prem covers text, images, audio and video in a single embedding space, running within a customer's environment and making no outbound calls once deployed. Small Jina models run on a single 8GB GPU while matching the accuracy of much larger models that need many times more GPU memory.

The suite includes all 28 Jina AI models, among them the jina-embeddings-v5-omni multimodal model and the jina-reranker-v3 reranking model. It installs with a single command and makes no licence-server, telemetry or model-registry calls. It supports both CPU and GPU hardware with automatic GPU detection, and applications reach it through standard API schemas so existing integrations work without rewriting code.

"Historically, teams running search and retrieval in regulated or disconnected environments have had to choose between capability and control," said Ajay Nair, general manager for Elasticsearch and Platform at Elastic. He said running Jina models fully on-premises removed that compromise.

Many Australian and New Zealand government agencies run OpenText Content Manager as their system of record while using Elastic as a cross-repository enterprise search layer over records, network drives and email. Jina On-Prem drops into that existing Elastic tier as a replacement for the models served through Elastic Inference Service, adding semantic, meaning-based retrieval across those repositories without re-architecting.

The decisive feature for that audience is data

sovereignty. Records classified at PROTECTED and above cannot be sent to external AI services, which has kept many agencies away from semantic search and retrieval-augmented generation. Running the models on-premises with no outbound calls lets an agency improve the findability of its records, and build AI assistants grounded in them, while staying inside its security boundary and IRAP obligations.

Beyond search, the same models underpin retrieval-augmented generation, where an AI assistant grounds its answers in an organisation's own documents. Running that pipeline on-premises means the grounding data never leaves the environment. Elastic said the small Jina models match the accuracy of far larger ones while running at a fraction of the compute cost, lowering the hardware barrier for agencies that cannot rely on cloud GPUs.

The release also lets air-gapped Elastic deployments adopt the models without changing how their applications call embedding and reranking services, easing the path for sites already invested in the Elastic stack. Jina On-Prem is available now for download via GitHub through an access token, with installation instructions on the Jina On-Prem Quick Start page.

www.elastic.co

Zuma Discovers and Governs AI Agents


Saviynt has launched Zuma, an identity security platform built to discover, govern and protect AI agents, large language models, AI-based applications and other non-human identities as organisations scale AI across the business.

The company said AI and non-human workers now outnumber humans by more than 100 to one, creating major risk that is currently unmanaged. Unlike dashboards that report risk after the fact, Zuma acts at runtime, evaluating an AI agent's intent, context and risk to stop unauthorised access in the moment.

Saviynt said it developed Zuma after working with large enterprises and finding that AI agents and non-human identities were multiplying across business functions, often without effective oversight. Legacy identity tools, it argued, cannot answer basic questions such as where all an organisation's AI agents are, who created them, what they are doing, who is overseeing them, and what control exists over the data and applications they reach.

Zuma answers those questions through three components in a single control plane. Zuma Insights continuously discovers AI agents, LLMs, AI-based applications and non-human identities, maps what they access and who owns them, and shows lifecycle changes on a timeline.

Zuma Access applies protection at runtime. Zuma Governance extends established identity best practices, such as ownership and access review, to AI.



"AI is fundamentally changing the identity security equation," said Sachin Nayyar, chief executive of Saviynt. "Enterprises are no longer securing only human users. They are now responsible for AI agents and non-human identities that can decide and access systems at machine speed."

Saviynt positioned Zuma as a single, end-to-end platform spanning visibility, runtime protection and governance, in contrast to what it described as fragmented point tools. Non-human identities, which include service accounts, machine identities, API keys and now AI agents, have become one of the fastest-growing categories in identity security. Their number and autonomy make traditional, human-centred review cycles hard to apply, which is the gap Saviynt is targeting with runtime, intent-aware controls that bring discovery, ownership and governance into one place.

<https://saviynt.com/>

Leidos Platform can Auto-Patch Software

Organisations can now autonomously find and fix software vulnerabilities using Parcata, a new platform from Leidos.

The model-agnostic tool harnesses a diverse set of large language models to accelerate vulnerability detection

in first and third-party software. Leidos said it can patch zero-day vulnerabilities in realtime, closing the gap between discovery and remediation that attackers exploit.

Being model-agnostic is central to the pitch. Rather than depending on a single frontier model, Parcata draws on several, which Leidos said ensures the mission is not tied to any one vendor's system. That design also lets the platform adapt as models improve or as particular models prove better suited to specific tasks.

The platform was developed by the company's Kudu Dynamics team, with technology validated through DARPA's AI Cyber Challenge, a United States programme that pitted automated systems against each other to find and fix flaws in real software.

It has demonstrated its capabilities through recent internal technical exercises and is being run against Leidos code before deployment to mission environments.

"The use of AI by cyber adversaries has changed the game. We must detect and remediate before a breach," said Jason O'Connor, president of Leidos Intelligence.

O'Connor said the tool would help organisations stay ahead of evolving threats and make sense out of the chaos, ensuring the mission was not dependent on any single frontier model.

www.leidos.com/

Microsoft Model Targets Vulnerability

Microsoft has launched MAI-Cyber-1-Flash, a cybersecurity AI model built to find vulnerabilities in complex codebases, inside its MDASH multi-agent identification and remediation harness.

The company said the combination delivers strong performance at half the cost of leading models. MAI-Cyber-1-Flash is designed to efficiently handle up to 90 per cent of tasks, letting MDASH reserve larger, costlier models for the hardest 10 per cent that truly need them. Microsoft framed the release as a response to a changing threat landscape. Attackers now use increasingly powerful AI to probe growing volumes of code for a single weakness. As the cost of finding a flaw falls, the company argued, the old model of scanning occasionally and patching eventually is obsolete.

MAI-Cyber-1-Flash was built to find challenging vulnerabilities in complex codebases and has been integrated into MDASH, which Microsoft said was honed by cybersecurity experts and hardened across a very large security estate. The company said the model beats rival systems, including Mythos, Gemini and GPT, on CyberGym, a benchmark for evaluating how systems reason over large codebases to find real vulnerabilities.

Microsoft reported the unified system of MDASH with MAI-Cyber-1-Flash scores 96 per cent on CyberGym, 12 points above the Mythos model. It said the combination delivers a 50 per cent cost saving against its current best MDASH configuration, which pairs several larger models. The company positioned the model as an example of a well-tuned, multi-model system, using a smaller specialised model for the bulk of tasks and escalating only the hardest cases to frontier models. It said this ensures the best model at the best price for each task.

Snowflake Adds AI Control Layer

Enterprises running AI agents across data, apps and platforms can now secure and monitor that activity through Cortex AI Gateway, a new control layer from Snowflake. The gateway lets organisations govern how first-party and third-party agents access models, tools, MCP servers and enterprise systems. It also intelligently routes agent requests and tracks token usage to help prevent runaway AI costs. Snowflake said the tool addresses two of the biggest barriers to enterprise AI adoption, securing agents and controlling AI consumption costs.

The launch reflects a shift the company describes as moving from data interoperability to agent interoperability. As AI agents increasingly collaborate across enterprise data, apps and platforms, they introduce security and governance risks that traditional architectures were not designed to handle. Many organisations also lack a central way to see and govern how much AI they are consuming.

Snowflake said Cortex AI Gateway governs both agents built within its platform, such as Snowflake CoWork and CoCo, and external agents developed on tools including Claude Code and Cursor. Teams gain a single point of visibility and control across AI access, agent activity and token usage.

The gateway builds on Snowflake’s May 2026 acquisition of Natoma, bringing that company’s enterprise Model Context Protocol platform directly into Snowflake for secure agent interoperability. MCP is the emerging open standard that lets AI agents connect to external data sources and tools, and its rapid adoption is a large part of why agent governance has become pressing.

Alongside the gateway, Snowflake announced secure third-party agent access integrations with Aembit, 1Password, Linx Security, Okta, SailPoint and Saviynt. The company said this set a new standard for how enterprises govern and audit AI agents across the security ecosystem.

www.snowflake.com

Cloud Platform for Engineering

Engineering teams in safety-critical industries can now manage controlled documents without local servers or VPNs, under a cloud-native platform launched by Synergis Software.

Synergis announced the general availability of Adept Cloud, a SaaS engineering document management platform. It is built for manufacturing, utilities, oil and gas, chemicals, pharmaceuticals and mining, where documentation accuracy affects operational safety, regulatory compliance and project outcomes.

Adept Cloud delivers the full capabilities of the company’s Adept platform in a browser-based environment that is hosted, maintained and continuously updated by Synergis. The vendor says the model removes the need for local infrastructure, VPN access or IT overhead.

The platform includes artificial intelligence capabilities built into the environment. Synergis Adept was named to G2’s 2026 Best Software Awards as a top 50 CAD and PLM product.

The launch extends engineering document control to organisations that manage large volumes of drawings and technical records but want to avoid running the supporting systems in-house.

“We built Adept Cloud because our customers told us where they needed to go — and we made a commitment to get them there,” said Kristen Tomasic, President, Synergis Software.

“What we built isn’t just a cloud version of Adept. It’s a modern, cloud-native platform - architected for the security, reliability, and scalability that asset-intensive organizations require for the decade ahead. Adept Cloud is also where we’ll continue to innovate - AI is built in, with more to come.”

<https://www.synergissoftware.com/>



Claude Powers Epiq eDiscovery Automation

Legal teams can now trigger complex eDiscovery tasks with plain-language requests under expanded capabilities in Epiq AI Accelerate , which turns natural-language queries into system-level actions across eDiscovery platforms.

The new capabilities let teams manage workspaces, act on data and surface insights while cutting the manual steps in time-consuming eDiscovery processes.

The capabilities are powered by Claude, the Anthropic AI model which interprets users’ natural-language requests. Epiq pairs Claude with a proprietary reasoning and execution layer that governs tool selection, reasoning and workflow execution across its platforms.

The workflows are accessible through the Epiq Service Cloud. They use Model Context Protocol functionality to act as a governed execution engine on top of Epiq Discover and other eDiscovery software. Built-in controls include user permissions, workflow approvals, data isolation, audit logs and AI action traceability.

Supported processes span workspace management, processing and imaging, document search, summarisation, review and export, and production set identification and conflict searches.

“The distinctive combination of Claude with the Epiq reasoning and execution layer is what allows Epiq AI Accelerate to support an unmatched and growing

set of capabilities,” said Roger Pilc, president, legal solutions at Epiq.

The Epiq Service Cloud has more than 2,600 clients and over 48,000 registered users. The new automation capabilities begin rolling out in August.

epiqglobal.com

Denodo Tackles AI Trust Gap

Denodo has released Denodo Platform 9.5, aimed at a problem stalling enterprise AI rollouts. AI agents can reach data but cannot consistently trust or understand it across a business.

The release adds two features of note. Metric views let organisations define core business metrics such as revenue, churn and customer value once, governed centrally. Dashboards, business intelligence tools and AI agents then work from the same numbers rather than recalculating them differently.

An expanded enterprise knowledge graph, delivered through the Denodo Data Marketplace, gives AI systems visibility into governance, lineage and business context around data, not just the data itself. Denodo describes this context layer as the increasing blocker for agentic AI deployments at scale.

Denodo Platform 9.5 expands its Data Marketplace with a new 360 graph and asset extensions, enabling organizations to define an enterprise knowledge graph with a broader set of data ecosystem assets and the relationships that connect them. Teams can

(Continued over)

define and manage related assets such as extract, transform, and load (ETL) processes, consuming applications, notebooks, business glossaries, data dictionaries, governance controls, data product contracts, data sharing agreements, AI skills, and other business or technical artifacts. The update also enhances reasoning in the Denodo Assistant to streamline data-view development, and expands connectivity across the data and AI ecosystem.

This release also expands connectivity across the modern data and AI ecosystem. These enhancements include improved support for unstructured data, new connectivity with Databricks and Azure AI Search, support for vector search with indexes for more efficient semantic and document queries, enhanced support for Delta tables and Databricks environments, and greater efficiency when using the Denodo Lakehouse Accelerator with Iceberg.

<https://www.denodo.com/en>

Xebia Prepares Data Estates for AI

A new platform combining proprietary AI agents with human data engineering teams promises to complete enterprise data migrations roughly three times faster than conventional human-led delivery, compressing projects that typically run 12 to 18 months into a few months. Xebia, a global digital transformation and engineering firm, announced general availability of Xebia Axis: Agentic Data Foundation. The solution assesses, builds, migrates, monitors and operates enterprise data platforms, positioning data estates for production AI workloads.

Six modules cover the data lifecycle. Readiness maps the data estate and produces a prioritised migration blueprint. Platform generates governed, policy-validated infrastructure as code.

Knowledge converts content held in documents, email and disconnected systems into a governed data layer for AI. Migration automates data and pipeline conversion with parity checks. Observability monitors data quality and agent activity, while Operations submits fixes as pull requests for human review.

Xebia says human teams set strategy and govern quality while agents execute with roughly 10 times the leverage of a purely human-led team. Access control, audit logging, data residency and bias controls are built into the platform, and data remains within the client's environment.

The platform works across Amazon Redshift, Google BigQuery, Databricks, Snowflake and Microsoft Fabric, and with large language models including Claude, OpenAI, Gemini and Mistral AI.

Organisations can start with a standalone readiness assessment, delivering a data readiness score, estate inventory and migration roadmap in two to four weeks. Assessment fees are credited toward a subsequent migration.

www.Xebia.AI/Axis

Whitevision Acquires UK's Documation

Intelligent document processing (IDP) vendor Whitevision has acquired Documation, a UK developer of purchase-to-pay and accounts payable automation software, extending its buy-and-build consolidation of the document automation market.

The deal, announced on 14 July 2026, is Whitevision's fifth acquisition since partnering with investor Main Capital Partners in August 2024, and its second in the UK after Document Logistix earlier this year.

Founded in 1993 and headquartered in Southampton, Documation specialises in automating financial document processing. Its suite covers the full downstream purchase-to-pay cycle, from purchase order creation and goods receipt to invoice matching, accounts payable processing and financial reporting. Customers span education, healthcare, manufacturing, corporates and municipalities.

"The combination further strengthens our Intelligent Document Processing capabilities and marks an expansion of our activities in the UK," said Whitevision CEO and founder Frank de Wit.

BeyondTrust Module Polices AI Agents

A new security module enforces what AI coworkers and autonomous agents are allowed to do on the endpoint, in real time, before they act. BeyondTrust's AI Agent Security applies uniform privileged-action enforcement across the AI tools widely deployed in enterprises, including Claude Code, Microsoft Copilot, Cursor and OpenAI Codex.

The module is the first of several planned for the company's Pathfinder platform. It is in private beta now, ahead of general availability in the second half of 2026, initially as an add-on to BeyondTrust's Endpoint Privilege Management product. Details are at beyondtrust.com. BeyondTrust says AI coworkers arrive on the endpoint carrying all the privilege of the user who launched them, at machine speed, without guardrails. From there they reach into SaaS, cloud and production systems, and have already erased production environments in some enterprises.

The module offers three capabilities. It discovers every AI assistant, copilot and autonomous agent running across endpoints, including shadow AI installed without IT approval, and maps what each can reach. It decides before the action, granting approved AI tools only the permissions they need with no default inheritance of full user credentials, and enforcing which MCP servers, plugins and external services agents may connect to. And it applies runtime controls that block unauthorised behaviours - such as credential exfiltration, deleting production code and database erasure - across Windows, macOS, Linux and containers, while keeping an audit-ready record.

<https://www.beyondtrust.com/>



AI-Native Output added to FineReader

ABBYY has released FineReader Engine 12.8.0 with support for DocLang, a new document format built for artificial intelligence rather than human readers.

DocLang gives developers a single, AI-readable structure for feeding documents to language models and agentic AI systems. ABBYY says the format cuts token usage and latency while improving structural accuracy. The standard was founded by ABBYY with IBM, HumanSignal, Nvidia and Red Hat, and sits under the Linux AI & Data Foundation. The working group argues that formats such as PDF, HTML and Markdown were designed for people.

That forces developers to build custom parsers at every integration point. The group says the patchwork raises compliance risk and increases the chance of model hallucinations.

DocLang aims to create a reliable layer between unstructured content and AI systems. ABBYY ran a controlled benchmark comparing the same document, model and task in PDF and DocLang form. The company claims the DocLang version improved output quality, lifted accuracy and reduced compute cost. An interactive version of the benchmark is published at [https://doclang-](https://doclang-benchmark.abbyy.tech)

[benchmark.abbyy.tech](https://doclang-benchmark.abbyy.tech).

The controlled experiment processes an annual report, a clinical study and vendor contract that represent the vast variety of enterprise documents rich with information created for human understanding but challenging for machines.

"ABBYY FineReader Engine is already used by thousands of organisations processing billions of documents every year," said Max Vermeir, vice-president of AI strategy at ABBYY.

"Now with DocLang as an AI native format, more companies will be able to accelerate innovation and have faster access to their business data to make smarter, more impactful decisions."

DocLang creates a reliable abstraction layer between unstructured data and intelligent AI systems. It standardizes the cacophony of digital document formats that enterprises operate on and gives AI systems the deterministic structure they need to perform reliably at enterprise scale.

Continued Vermeir, "DocLang is specifically engineered to address industry challenges with a minimal, standardized, and AI-native method for representing document structure, meaning, layout, and governance."

<https://www.abbyy.com>

Data Governance for Process Records

Manufacturers in regulated industries can now capture process data at the point of measurement to create secure, time-aligned digital records for compliance and audit, with the launch of Watlow's Edge Process Management (EPM) platform. The first release, EPM Data, targets environments where process records must withstand regulatory inspection and customer audit.

The platform combines data management, process control and automation for temperature-driven operations in a modular, distributed architecture. It is designed to replace fragmented systems and legacy chart recorders across industries including metals, heat treatment, life sciences, power generation, water and wastewater, and food and beverage.

"The process record has effectively become part of the product," said Chris Mooney, Industrial Controls Segment Director at Watlow. "With EPM Data, and the broader platform behind it, we are giving customers a trusted, time-stamped record created at the point of measurement, designed for daily operations and audit readiness, without adding system complexity."

Capabilities include high-precision inputs with field calibration accuracy for standards such as AMS2750 and CQI-9, and secure, tamper-resistant records designed to support requirements including FDA 21 CFR Part 11 and Annex 11. A single consistent process record can be used by operations, quality and audit teams for review, approval and long-term retention.

The initial release includes two panel-mounted devices - the compact EPM-07 for smaller furnaces, ovens and skids, and the EPM-15 for larger or multizone systems - plus Data Reviewer software for visualising historical process data, with optional audit trail and electronic signature functionality.

EPM Data is being introduced to selected customers, with broader availability planned through Watlow's global sales network.

www.watlow.com/epm

Law Firm Data Sprawl in Sights

Law firms across Australia, New Zealand and the wider APAC region gain a combined consulting and technology path for finding and governing client data that has leaked beyond the document management system, under a new strategic partnership announced for the legal sector.

The alliance pairs Verlat Consulting, a legal technology consultancy with offices across Australia, New Zealand, Singapore, the United Kingdom, Japan and Malaysia, with ActiveNav, a data discovery and information governance vendor. It combines Verlat's recently established information

governance practice with ActiveNav Cloud, a platform built to discover unstructured data at scale.

The partnership targets a problem the companies say plagues firms of every size: matter documents that have spread beyond the DMS into network file shares, cloud storage, local drives and other unstructured repositories.

ActiveNav Cloud discovers and classifies that content, then enables firms to remediate it - importing documents into NetDocuments for proper governance, disposing of them under retention policies, or securing them in place.

"Law firms invest heavily in document management systems, but the reality is that a significant volume of sensitive client data still ends up outside those systems ungoverned and often invisible," said Jason Mills, CEO of Verlat Consulting. "With ActiveNav, we can finally give our clients a clear picture of that exposure and a practical path to resolving it."

The companies also position the platform's outputs - clean, curated datasets - as groundwork for law firm AI initiatives, alongside retention enforcement, cleanups and migrations.

EQS Embeds AI In Compliance Platform


Compliance teams using EQS Group's platform now have an AI layer, Q by EQS, operating directly inside the system of record with access to case data, disclosures, policy libraries, due diligence results and workflow history. AI-powered whistleblowing triage, case classification and policy assistance are available now, with agentic capabilities to follow later this year.

The launch draws on the EQS AI Benchmark Report Volume 2, which the company describes as the industry's first evaluation of AI performance on compliance work. It found frontier AI models achieve up to 87 per cent accuracy on real-world compliance tasks, with the top four models within a single percentage point of each other - making the surrounding expertise, workflows and governance the real differentiator.

Initial use cases include identifying involved parties in whistleblowing reports and automatically blocking their case access, report triage and severity assessment, investigation support including evidence summarisation, disclosure risk scoring, and natural language data queries with opt-in benchmarking against anonymised EQS peer averages.

"Building AI that works in compliance is not a model problem - it's a domain problem," said Moritz Homann, Head of AI at EQS. "That includes human oversight at important decision points and full auditability."

The planned agentic layer will execute multi-step tasks in the EQS Compliance Cockpit under human oversight. Compliance teams can configure where Q assists, where it is restricted and where it is disabled, and customer data is never used to train AI models operating in EQS products. Details are at eqs.com.



Agentic Gateway Targets Compliance

Enterprises in regulated industries can now route AI agents through a central orchestration layer that manages third-party MCPs for compliance and reduces token consumption, under major updates to Perforce Intelligence announced by Perforce Software. The release spans three components: an agentic gateway, an AI-assisted testing platform driven by natural language, and a unified compliance layer that turns written security policies into continuous enforcement.

The Perforce Agentic Gateway, available now through GitHub, provides a centralised enterprise access layer for the Perforce MCP portfolio through a single install. It can also manage third-party MCPs, enabling organisations to govern AI where software is built and delivered - across code, IP, data, infrastructure and testing.

Perforce Autonomous Testing lets non-testers describe what they want to validate in natural language through a single chat interface, with AI executing across functional, performance, mobile, web, desktop and accessibility testing. A single test can work across desktop web, iOS and Android. The platform is built on Perforce's existing BlazeMeter

and Perfecto tools, with a Delphix integration planned for test data provisioning.

The Unified Compliance platform sits above existing Puppet deployments. It translates internal and external policies into code, routes enforcement to the applicable infrastructure including Kubernetes environments, monitors for drift, remediates violations and maintains audit evidence across on-premises, hybrid and multi-cloud environments.

"There's a lot of money being spent with AI, but the next evolution needs to be on ROI of AI, which can only be measured if you have control mechanisms in place and visibility into what it's doing," said Jim Mercer, Program Vice President at IDC, commenting on the release.

Anjali Arora, CTO of Perforce, said the company aims to be the control plane for enterprise AI workflows. "The future of AI in the enterprise will be defined by how effectively organizations can control, scale, and operationalize AI workflows, models, and deliver real business results with clear visibility and traceability," she said.

A 2026 Harris Poll survey, commissioned by Dataiku, of more than 900 CEOs, found 80 per cent say their role will be at risk if their company fails to deliver measurable business gains from AI by the end of 2026.

<https://www.perforce.com/>

NetDocuments Unveils Context Graph

Lawyers opening an unfamiliar matter will see its full context - summary, key parties, activity timeline, relevant precedent and the people who have done the work before - under a reimagined NetDocuments platform built around what the company calls the first context graph for legal work.

AI agents working inside the platform, or in external tools connected via MCP including Claude and ChatGPT, draw on the firm's institutional knowledge rather than a single session's uploads.

NetDocuments says the system continuously maps how every matter, document and communication in a firm connects, across hundreds of millions of records, while respecting existing permissions and ethical walls.

The company built the system with partners including AWS and Elastic to continuously process and connect documents under strict governance.

For AI applications, the shift addresses a structural limit: until now, legal AI has operated on whatever a lawyer could upload into a single session.

Platform capabilities include search by meaning across the full firm with permissions preserved, automatic matter overviews, and AI-powered drafting, redlining, playbook generation and document comparison across NetDocuments and Microsoft 365.

The new platform launches alongside the original NetDocuments interface.

Both experiences are fully supported, share the same underlying data and operate under the same governance model, allowing firms to move on their own timeline.

<https://www.netdocuments.com/>

A Record-Keeping Tool for AI Agents

CellTrust has added a feature to its SL2 Enterprise platform that captures and archives communications involving AI agents, extending compliance record-keeping to agentic systems.

The new SL2 AI Capture records interactions across human-to-agent, agent-to-human, agent-to-agent, agent-to-system and system-to-agent exchanges.

The company says this keeps agent communications under existing regulatory frameworks.

CellTrust argues the capability also makes it easier for third-party surveillance tools to flag deepfakes and bots.

The platform feeds Digital Communications Governance and Archiving systems. It unifies channels including SMS, WhatsApp, Microsoft Teams and voice.

The launch responds to fast-growing use of multi-agent AI in regulated sectors such as financial services, government, legal and healthcare.

CellTrust says these deployments create productivity gains alongside a widening compliance gap.

The company says SL2 AI Capture delivers capture at the point of origin, end-to-end traceability, immutable archival and cryptographic verification.

It frames this as a provable chain of custody rather than the basic audit logs many AI platforms provide.

<https://www.celltrust.com>

Mistral OCR 4 Targets RAG/Search

French developer Mistral AI has released OCR 4, a document-recognition model that returns bounding boxes, block classification and confidence scores alongside extracted text.

The model moves beyond clean text and tables to a structured view of each page. It labels blocks such as titles, tables, equations and signatures, and scores its confidence per page and per word.

Mistral says that structure suits retrieval-augmented generation, agent workflows and compliance tasks such as redaction and human-in-the-loop verification. The output feeds citation-ready inputs for search and ingestion pipelines.

OCR 4 is an ingestion component of Search Toolkit, Mistral's open-source, composable search framework.

OCR 4 accepts common enterprise formats, including PDF, DOC, PPT, and OpenDocument and supports 170 languages across 10 language groups.

It runs in a single container for fully self-hosted deployment. That lets organisations keep document data inside their own infrastructure for residency and compliance.

Mistral claims independent annotators preferred OCR 4 over every rival system tested, with average win rates of 72 per cent.

The company says OCR 4 also leads the public OlmOCRBench.


Mistral says rival scores reflect its own internal reproductions and recommends testing on your own documents. It also flags limits in common benchmarks, noting that errors in reference data and equivalent maths notation can mark correct output as wrong.

Through the API, OCR 4 is priced at US\$4 per 1,000 pages, halving to US\$2 with a batch discount.

The model is available via Mistral Studio, Amazon SageMaker and Microsoft Foundry, with Snowflake support listed as coming soon.

Mistral notes OCR 4 is a document-understanding model, not a decision-maker. It is not intended for medical, legal or safety-critical use.

<https://mistral.ai>



Action Targets Metadata Drudgery

Manual metadata upkeep, a long-standing bottleneck for data governance teams, is the target of a new AI agent from Actian, the data and AI division of HCLSoftware. The Actian Data Steward Agent is embedded in the Actian Data Intelligence Platform and is available now. It automates documentation, enrichment, classification and governance work across the data catalogue.

The agent writes and updates asset documentation as data pipelines evolve. It assigns ownership to orphaned assets and flags gaps in documentation. It also recommends classifications, including personally identifiable information (PII) and sensitivity levels, and maintains enterprise-wide glossary definitions and business terms.

Action claims the agent covers the full scope of the data steward role, rather than automating isolated tasks in the manner of AI copilots. Suggestions are grounded in lineage, usage patterns and existing catalogue context. Human stewards retain final approval, reviewing and validating outputs rather than creating them from scratch.

The agent is built on the platform's federated knowledge graph, semantic layer, data products and data contracts. External AI agents connected via Model Context Protocol (MCP) or agent-to-agent (A2A) integrations can operate within the same governed business context.

<https://www.actian.com>

Templafy Tackles Compliance Gap

Templafy MCP is a new integration that connects third-party AI platforms with Templafy's document agents. The company says this lets enterprises govern AI-generated document creation across teams, tools and models without forcing employees into a single AI interface. The Model Context Protocol (MCP) integration allows AI tools such as ChatGPT, Claude, Copilot, Gemini and Perplexity to hand content to Templafy's document agents. The agents apply company-approved templates, brand assets, prompts, formatting rules and business data. The result is editable documents in Microsoft 365 formats that meet organisational standards.

The launch addresses a widening compliance gap. Generative AI has accelerated content creation across enterprises, but drafts routinely require manual rework to meet accuracy, brand and compliance standards before they can be shared. Documents produced outside approved templates also create records management and version control headaches.

MCP has rapidly emerged as the standard mechanism for connecting AI models to enterprise tools and data sources. Templafy claims its MCP is purpose-built for document generation rather than general tool connectivity. The document agents also draw on enterprise systems such as CRM and digital asset management platform.

<https://www.templafy.com>

Data Protection for AI Agents

Enterprise AI platforms can now reach data protection and cyber resilience operations through a new open-standard interface. Cohesity has launched Cohesity Maestro to make its data cloud natively accessible inside existing corporate AI tools.

The system uses the Model Context Protocol to connect with Anthropic Claude, OpenAI ChatGPT and Google Gemini. This approach removes the need for custom integrations or proprietary connectors.

Cohesity said the launch delivers the industry's first headless data security architecture. Operations teams can trigger restores, hunt threats and orchestrate recovery without switching to a separate console. Actions run through autonomous agents inside the AI workflows teams already use.

The architecture stays governed by existing role-based access controls, authentication and audit frameworks.

This keeps data insights under the same security terms that apply to direct platform access. Teams can also ask what changed across their environment and surface a prioritised view of critical risks.

The interface exposes core capabilities including cyber resilience orchestration, realtime telemetry and threat signals. It also connects Cohesity Gaia, an AI search engine offering semantically enriched queries over protected data.

Cohesity Copilot and RecoveryAgent are available now, while the Maestro interface is expected later this year.

<https://www.cohesity.com>

Project Hourglass Secures AI Workflows

A new alliance of global systems integrators will deliver dedicated security and resilience architectures for autonomous software agents. Rubrik has partnered with Cognizant, Deloitte, LTM, HCLTech, NTT Data and Wipro under an initiative called Project Hourglass.

The integrators are building Rubrik Agent Cloud for Anthropic's Claude Code into their cybersecurity and digital transformation work.

Organisations face mounting risk as autonomous tools write and deploy code faster than security controls can adapt. Rubrik Zero Labs data shows 86 per cent of cybersecurity leaders expect AI agents to outpace their guardrails within a year.

The platform provides behavioural guardrails, observability and a control-plane resilience system that backs up the rules governing agent behaviour.

Rubrik said the platform offers the industry's only agent rewind, reversing unintended automated actions. It also provides immutable codebase

recovery when a destructive action bypasses standard version control. The technology combines three layers: Runtime Agent Security, Agent Rewind for repository recovery, and AI Context Guard for prompt integrity.

Project Hourglass forms part of Rubrik's Agentic Resilience Partner Program. The program gives integrators joint go-to-market support, technical certification tracks and access to Rubrik's engineering teams.

Partners are embedding the system into their own platforms, including Cognizant's Neuro AI and HCLTech's VERITY Frontier AI Resilience framework. Rubrik Agent Cloud for Claude Code is generally available.

<https://www.rubrik.com>

Compliance Platform Verifies AI

Enterprises can now manage autonomous software sprawl and rising operating costs through a single governance platform. Thoughtworks has launched Agent/works, giving organisations one control plane and a governed runtime deployable across any cloud.

The platform targets a sharp shift from early AI experimentation to high-stakes operational deployment.

Thoughtworks unveiled it at the Databricks Data + AI Summit in San Francisco, partnering with Databricks on enterprise AI governance.

According to Sonar's 2026 State of Code Developer Survey, 42 per cent of committed code is now AI-generated or AI-assisted. Separate research from AppSec Santa found 25 per cent of AI-generated code samples contained critical security vulnerabilities.

Unmonitored autonomous systems risk exposing sensitive data, breaching compliance rules and obscuring fast-growing AI operating costs.

Agent/works analyses every path through a workflow and confirms a fully compliant route exists before an agent runs.

It grants capability-based, scope-bound and time-limited permissions that narrow automatically as agents touch sensitive systems.

The governed runtime applies adaptive policies during execution, covering both end-to-end workflow agents and interactive coding tools.

A centralised registry provides visibility, evaluations, usage analytics and cost controls across every agent, model, tool and policy.

The platform runs on a multi-model backend, registering any model through a standard API and connecting third-party tools and services.

This architecture lets enterprises scale automation safely without governance becoming a bottleneck.

<https://www.thoughtworks.com>



Clario Targets Stale Files Undermining AI

US startup Clario has launched from stealth with \$US6 million in seed funding, targeting the redundant, obsolete and trivial files that clog enterprise systems.

The company connects to platforms including Google Drive, SharePoint, OneDrive, Box and Confluence. It scans metadata to surface low-value data, then prompts staff through Slack or Teams to keep, archive or delete each file.

Clario says it charges only when a customer makes one of those decisions, and that the system learns from them over time to automate cleanup.

Over time, Clario learns from these decisions to build an autonomous system that routinely cleans, organizes, and maintains unstructured data.

Clario cites industry estimates that 78 per cent of stored enterprise data is unstructured, with more than a third effectively garbage.

It also points to a Gartner forecast that 60 per cent of AI projects will be abandoned this year due to poor data quality.

The company claims to be the first product purpose-built to eliminate data ROT, though established data-governance and storage vendors address adjacent problems. Early customer scans surfaced terabytes of waste, including obsolete file formats and feature-length films left by former staff.

<https://clarioclean.com>

Data Prep Drives Microsoft Deal

Data preparation, one of the biggest barriers to moving AI from pilot to production, is the focus of an expanded collaboration between Unstructured and Microsoft. Under the arrangement, Unstructured's cloud-native ETL platform is integrated with Microsoft Azure services to prepare complex enterprise content for AI workloads. The platform parses, chunks and enriches data from more than 64 file types. The output feeds large language models, retrieval-augmented generation (RAG) pipelines, copilots and agentic workflows.

Native connectors ingest data from Azure services such as Azure Blob Storage. The prepared data can then be indexed in Azure AI Search and used with Microsoft Foundry to support production AI applications.

The collaboration targets highly regulated industries, including financial services, healthcare, insurance, pharmaceuticals and government. Unstructured can run within customer Azure environments, allowing organisations to retain their own security, compliance and data governance controls rather than sending content to an external service.

Enterprises can also procure the platform through Microsoft Marketplace, aligning purchases with existing Azure spending commitments. Unstructured supports more than 30 content connectors, including Microsoft OneDrive, SharePoint and Azure Blob Storage.

<https://unstructured.io>

Document Workflows Expand at MamaSign

Businesses that send client-facing agreements can now run the entire signing journey under their own brand, from signing links and emails to completed PDFs, following an expansion of MamaSign's whitelabel eSign and document workflow platform. The update adds custom-domain document links, custom SMTP email sending, API access and webhooks.

The platform is built around a document journey: prepare a file, send it for electronic signature through a branded link, deliver related invoice or proposal pages, verify the completed PDF and connect document events back to business systems.

A signed PDF verification workflow uses document identifiers and SHA-256 fingerprint comparison, a cryptographic file-matching method. It gives recipients a way to check whether a completed file is recognised by the platform after it has moved through email, portals or cloud storage.

The platform targets small businesses, agencies, consultants, HR teams and service providers that regularly send contracts, approval documents, invoices, proposals and onboarding paperwork.

<https://mamasign.com/>

Inbox Triage for Malicious Emails

Barracuda Networks has launched Integrated Email Protection, a cloud email security product that detects and removes threats even after a message reaches the inbox. The product layers on top of Microsoft 365 and Google Workspace. Barracuda says it continuously re-evaluates messages and can claw back malicious email after delivery, rather than scanning only at the point of entry.

It runs on the BarracudaONE platform and correlates signals across email, identity, network, data and applications. The company says this cross-domain view helps managed service providers investigate and remove risk faster.

Barracuda released research alongside the launch. Its red team recreated a multi-stage AI-powered attack that moved from one phishing email to identity theft, MFA bypass and endpoint compromise in minutes. The company says one in seven compromised accounts is now used to launch further attacks, a figure it expects to rise as attackers automate with AI.

New features include AI agents that triage and remediate threats tenant-wide, a unified quarantine that consolidates Microsoft-quarantined email into Barracuda, and value reporting that quantifies threats stopped. The Bailey AI assistant explains each verdict in plain language and lets teams reverse actions.

<https://www.barracuda.com>

Document Lifecycle within Salesforce

Documents for Salesforce is a new natively built application from Nutrient that runs document generation, editing, redaction and signing inside the CRM. The company says document workflows are often the weakest link for organisations that anchor revenue, legal and compliance work in Salesforce. Quotes, contracts and approvals are typically generated, edited and signed in disconnected tools, then returned to shared drives.

Nutrient Documents for Salesforce keeps templates, generation, editing, redaction, signing and envelope tracking on the Salesforce record. The workflows run under the permissions and controls organisations already operate, the company said.

Administrators build templates without writing queries, using merge fields, conditional logic and related-record references. Generation pulls live record data from standard or custom Salesforce objects, which the company says removes snapshot drift between the CRM and the document.

Users review and edit PDF and DOCX content in a Salesforce-native editor, with markup and comments tied to the record. Sensitive content can be redacted before sharing, and signing is initiated from the record.

The company says the approach cuts tool sprawl and tightens audit posture, with version and approval history kept on the record. It aligns with SOC 2 Type II compliance. Nutrient also positions the platform as groundwork for AI agents taking on document work.

<https://nutrient.io>

Commvault Speeds Clean Recovery

Commvault has integrated its Threat Scan capability with Google Threat Intelligence. The move aims to help organisations identify clean recovery points faster and cut the delay between a cyberattack and a safe restore.

The integration incorporates Google Threat Intelligence data and scanning into Commvault Threat Scan workflows. It addresses a common problem after an attack. Security teams may quickly identify indicators of compromise, but recovery teams still need to validate backup data before a restore can begin.

Google Threat Intelligence combines Mandiant frontline intelligence, VirusTotal's crowdsourced data and Google's own threat insights. Integrating it with Threat Scan lets customers analyse protected workloads for malware and pinpoint which recovery points are compromised. Threat context is provided for further investigation.

As part of the release, Commvault is adding scanning that collects file hashes inline during backup operations. The company likens file hashes



to fingerprints. They give teams a fast way to check recovery points against threat intelligence indicators and identify clean files for recovery.

The inline inspection lets customers start with rapid threat intelligence validation, then selectively run deeper malware, encryption and forensic analysis when needed.

Commvault says the layered approach speeds recovery decisions while maintaining confidence in restored data. The capabilities also strengthen its Synthetic Recovery feature, which detects and removes threats during recovery while keeping clean data intact.

<https://www.commvault.com>

Cloudflare targets AI agents

Cloudflare has expanded its Agent Cloud platform with a set of new tools aimed at moving AI agents from development experiments to production-grade enterprise workloads.

The expansion includes Dynamic Workers, a new isolate-based runtime for running AI-generated code in sandboxed environments. Cloudflare says Dynamic Workers spin up in milliseconds and cost significantly less than traditional containers. The

company claims they run at 100 times the speed of containers at a fraction of the cost.

Cloudflare is also introducing Artifacts, a Git-compatible storage layer built for AI-generated code.

The company says Artifacts allows developers to create tens of millions of repositories and fork from remote sources, giving agents persistent storage accessible through standard Git clients.

Sandboxes, a persistent Linux environment for agents requiring a full operating system, has moved to general availability.

A sandbox provides a shell, filesystem, and background processes, enabling agents to install packages, run builds, and iterate across tasks.

A new SDK framework called Think is designed to support long-running, multi-step agent tasks, addressing what Cloudflare describes as the limitation of agents that can only respond to single prompts without persisting state across sessions.

Cloudflare is also expanding its AI model catalogue, following its acquisition of Replicate, to include proprietary models from OpenAI and open-source alternatives through a single interface.

The company says developers can switch between model providers by changing a single line of code.

<https://www.cloudflare.com>