idm.
information & data manager
AUGUST-SEPTEMBER 2026

Terabit Attack Exposes Botnet Blind Spot

Document Format Pitched as PDF for the AI Era

Why Sovereign Cloud is Harder than it Sounds

Cyber Resilience Foundations Still Matter in the AI Era

Data Sovereignty Drives Hyland APAC Push

Hyland has significantly widened its presence in Australia and the Asia-Pacific region, bringing new cloud infrastructure, major cloud partnerships, and a raft of AI platform capabilities aimed at organisations managing complex, document-driven compliance workflows.

The company has launched its Content Innovation Cloud on Amazon Web Services' Asia Pacific (Sydney) region, addressing a central concern for regulated Australian industries: data residency. Building on an existing strategic collaboration agreement with AWS, the deployment is aimed at sectors including healthcare, financial services, insurance, and government, where keeping data within Australian borders is a compliance requirement rather than a preference.

"With demand for AI-driven enterprise solutions accelerating globally, organizations across the Asia-Pacific region want to take advantage of AI but are mindful of data sovereignty and deploying advanced content and AI capabilities within their own regional infrastructure," said Tim McIntire, chief technology officer at Hyland.

"By extending the Content Innovation Cloud into Australia, Hyland is removing that barrier, empowering our customers to move forward with confidence as they modernize core operations and adopt the agentic enterprise."

Simultaneously, Hyland announced a separate strategic partnership with Microsoft to make the Content Innovation Cloud available on Microsoft Azure globally. The deal includes a co-sell and go-to-market arrangement, and will see Hyland solutions listed on the Microsoft Marketplace - a commercial step that allows enterprise customers to use existing Azure spending commitments to procure Hyland products, reducing procurement friction that has historically slowed platform consolidation in large organisations.

"Our customers are at different stages of their cloud and AI journeys, and this partnership ensures Hyland can support them wherever they are, while building toward a shared future powered by intelligent, agent-driven work," added McIntire.

Underpinning both cloud moves, the company announced general availability of the Enterprise Context Engine, which uses industry-specific ontologies to give AI systems domain-aware context when processing enterprise content. In practice, this means an AI agent working on a health record can understand the relationships between diagnoses, medications, and clinical notes - rather than treating them as unconnected documents.

Also reaching general availability is the Enterprise Agent Mesh, a layer for coordinating multiple AI agents operating across an enterprise. Hyland introduced a companion Control Tower capability providing realtime performance metrics, the ability to pause or adjust agents based on defined thresholds, and a governance framework for approving agents before deployment. The company described these as tools for organisations to actively manage rather than simply monitor their AI deployments - a distinction that carries practical weight as AI agent proliferation creates new operational and audit risks.

A new Agent Lifecycle Management framework extends governance across the full lifespan of an AI agent. Components include an Agent Passport - a standardised record of an agent's identity, capabilities, and compliance status required before production deployment - and an Agent Library providing a searchable catalogue of all agents in an organisation's ecosystem, tracking ownership, function, and version history.

Hyland also introduced a headless mode for the Content Innovation Cloud, exposing core content, context, and governance capabilities as APIs. The move enables data engineering teams and independent software vendors to embed Hyland's content intelligence into their own applications and third-party AI tools without adopting Hyland's own interface. The company cited integration with platforms including Databricks and Snowflake as target use cases.

The product announcements are accompanied by preconfigured industry solutions. Hyland described an Agentic Hospital solution it claimed could enable up to five times faster referral assembly and 60 per cent more efficient record intake.

An Agentic Accounts Payable solution is said to target up to ten times faster invoice cycle times and a 60 per cent reduction in cost per invoice. An Agentic Bank solution was presented as reducing the time for loan applications to reach underwriter-ready status from weeks to days.

Hyland will be bringing its annual user conference, CommunityLIVE, to Melbourne on August 25, 2026.

https://www.hyland.com/en

idm
information & data manager

Publisher/Editor: Bill Dawes
Email: bill@idm.net.au

Web Development & Maintenance: Cordelta

Advertising Phone: 02 90432943
Email: idm@idm.net.au

Published by Transmit Media Pty Ltd
PO Box 392, Paddington NSW 2021, Australia

New Document Format Pitched as PDF for the AI Era

A new open document format backed by IBM, NVIDIA, and Red Hat is being positioned as the AI-era successor to PDF. The LF AI & Data Foundation announced the formation of theDocLang Specification Working Group on June 9.

DocLang is an open, universal, AI-native format designed to standardise how enterprises prepare, exchange, and govern document data for AI systems. Its backers want it to play the role for machines that PDF has played for human readers.

“Our vision is for DocLang to become a broadly adopted international standard for AI-ready documents, providing a consistent representation for both humans and machines, much as PDF became the universal standard for document exchange in the human-centric era,” said Peter Staar, Principal Research Scientist and Manager at IBM Software.

The working group operates under the Linux Foundation's LF AI & Data body. Founding members IBM, NVIDIA, and Red Hat are joined by contributors ABBYY and HumanSignal. It will run under the Joint Development Foundation's vendor-neutral, open governance model. Its goal is a specification supporting reliable, interoperable document processing across AI and agentic workflows.

A key feature for compliance teams is embedded governance controls. These help downstream systems enforce policies on privacy, extraction scope, and model training permissions.

The specification will also preserve semantic meaning and geometric layout in a single format. It encodes structural elements such as headings, paragraphs, and tables alongside their position on the page. The format is optimised for modern AI tokenisation and modelling approaches.

Enterprises currently work across a fragmented landscape of formats, including PDFs and JPEGs, built for human rather than AI consumption.

The foundation said this disconnect can introduce complexity, raise costs, and reduce reliability when extracting meaning from business documents.

“Documents remain one of the most important sources of enterprise knowledge, but most were never designed for AI-driven workflows,” said Mark Collier, general manager of AI & Infrastructure at the Linux Foundation and executive director of LF AI & Data.

“With the launch of the DocLang Working Group, we are bringing the open source community together to develop a vendor-neutral, interoperable standard that helps organizations prepare document data for AI more reliably, transparently, and at scale,” Collier said.

Staar said DocLang draws on years of IBM research into document representation for AI, including OTSL for compact tables and DocTags for preserving structure. ‘

“Together with our industry partners, we have distilled these lessons into DocLang, a new AI-native format for unstructured content,” he said.

Maxime Vermeir, Vice President, AI Strategy at ABBYY, said the format addresses a foundational problem in enterprise AI. “Documents were built for humans, not machines,” he said. “By introducing a minimal, standardized, and AI-native representation of document structure, layout, meaning and governance, DocLang creates a far more deterministic foundation for modern AI systems,” Vermeir said.

The working group builds on Docling, the open source document processing toolkit hosted by LF AI & Data. Docling was developed by the AI for Knowledge team at IBM Research Zurich and released as open source in 2024.

Docling converts formats including PDF, Word, PowerPoint, Excel, HTML, and images into structured, AI-ready outputs. DocLang complements it by defining an open standard for expressing and exchanging that structured output across systems. Together, the two projects span document ingestion, parsing, standardised representation, and downstream consumption by language models and agents.

https://doclang.ai

Federal FOI Administration Fails Pro-Disclosure Test

The Albanese government's pre-election commitments to open government and transparency appear not to have made it through to key federal portfolios, which are failing the pro-disclosure obligations of the Freedom of Information Act according to the Australian National Audit Office (ANAO).

Three senior Commonwealth departments were examined for the report: The departments of Prime Minister and Cabinet; Treasury; and Infrastructure, Transport, Regional Development, Communications, Sport and the Arts. The scope of report was 1,111 FOI requests seeking access to documents other than those containing the applicants own personal information, which were active at any time during 2024 across the three agencies.

The (ANAO) found that the Department of the Prime Minister and Cabinet (PM&C), Treasury and Infrastructure released documents in just 43 per cent of sampled FOI requests. The report concluded administration of FOI Act requests by the audited entities has been only "partly effective in giving the community access to information".

"It is not evident that entities are meeting their obligation to take reasonable steps to assist applicants make FOI requests. None of the internal guidance of the entities outlined what the entities considered to be reasonable steps to assist applicants. This approach is not prodisclosure," the report concludes.

The audit also identified a pattern of "courtesy consultations" with ministers' offices not provided for in the FOI Act.

Infrastructure consulted its Minister's office on 82 of 197 sampled requests. PM&C consulted on 42 of 210.

The ANAO found a "lack of records in all three audited entities to be able to determine whether advice or feedback led to a change in decision by the decision-maker".

Poor records management was identified as a key driver of compliance failures across the three departments.

The ANAO found that "The three audited entities are unable to demonstrate that they have met the obligation to take all reasonable steps to find requested documents Entities are not maintaining adequate records of the searches they undertake in response to FOI applications to assure completeness. Entity records also do not consistently support decisions to refuse access because records could not be located or do not exist."

The Department of the Prime Minister and Cabinet acknowledged that record-keeping was an area for improvement and outlined that it undertook a 'new weekly process to remind staff to save records appropriately' including:

'a weekly calendar appointment to schedule and remind staff to complete record keeping. This is in addition to a one-off records management day where a senior adviser paid for lunch for the staff who spent the entire day saving records for FOI requests to the official record repository; and

'create and maintain SOPs [standard operating procedures] for repeated aspects of the FOI process'.

The Department of the Treasury advised ANAO in December 2025 that it:

- 'has instituted regular filing time twice a week, and staff are regularly reminded to file their FOI matters completely and promptly'; and
- is 'also considering how we can make our record keeping more robust, including by investigating options for dealing with technical barriers to easier record management'.

Treasury issued its first FOI policy on 17 February 2026 after the audit work concluded. It had no finalised FOI policy or procedures in place at the time of the audit

Treasury's draft policy directed staff to send FOI decisions on Fridays. The ANAO found this delayed releases and breached the Act's "as soon as practicable" requirement.

PM&C's "Business Rules" document had no executive endorsement, no version control, no release date and made no reference to the objects of the FOI Act.

Infrastructure's June 2024 policy was the strongest of the three, with version control, formal endorsement and reference to obligations under the Archives Act 1983.

For 79 per cent of FOI requests examined by the ANO, the three agencies had not met their own requirements to evidence the reasonableness of document searches.

A record of the decision was not evident for 30 per cent of FOI requests examined. The record of the decision provided to the applicant was not retained for 11 per cent of requests. In 47 per cent of sampled requests, information in the case management system could not be verified to records.

The OAIC reported in 2024-25 that 43,456 FOI requests were received by agencies and ministers, a 25 per cent increase on the prior year. Of 25,211 decisions made in 2024-25, 79 per cent were to refuse access in full or in part. In the same period, 62 per cent of Information Commissioner reviews set aside the original decisions of agencies.

The OAIC has said this "high set-aside rate indicates there is more work to do to embed a pro-disclosure approach".

Statutory timeframes were not consistently met across the three departments. For 63 per cent of requests examined, the entities used statutory extensions beyond the 30-day decision period. Even after extensions, 29 per cent of requests with adequate records were processed outside the statutory deadline.

Entity disclosure logs were also incomplete, with 8 per cent of expected disclosures not made and 23 per cent of published disclosures late. Documents were exempted or information was redacted in 53 per cent of the sampled FOI requests. Decisions to exempt a document or redact information were "often not recorded, beyond the decision notice to the applicant", the audit found. Entity records "do not demonstrate a pro-disclosure approach to FOI requests", according to the report.

The full report is available here.

Frozen in COVID: Victoria Police's nine-month FOI delays

Victoria Police is taking an average of nine months to process Freedom of Information requests – unchanged since at least 2022, when it was attributed to COVID-19.

The force's FOI page at police.vic.gov.au continues to advise: "We are currently experiencing a significant increase in Freedom of Information (FOI) submissions. It will take 35 weeks on average to process your FOI request."

The delay stands in stark contrast to the FOI Act's requirement that agencies respond within 30 to 45 days.

Victoria Police told ABC News that a new records management system would be implemented "soon".

In May 2024, RecordPoint announced it had been awarded a seven-year, $A7.2 million contract to replace the Law Enforcement Assistance Program (LEAP) - Victoria Police's electronic document and records management system used to classify, store, access and manage documentation across the force's 23,000-plus staff.

LEAP has been a long-running problem for the force. A 2016 report revealed Victoria Police had already extended LEAP's life by seven years after a planned replacement project was abandoned in 2011, having cost $A45 million without delivering a workable outcome.

OVIC data shows record demand, worsening compliance

The Victoria Police delays sit within a broader picture of FOI demand surging across Victoria. The Office of the Victorian Information Commissioner (OVIC) 2024-25 annual report recorded 58,181 FOI requests across Victoria in 2024-25 - a 9.4 per cent increase from the previous year and another record high.

Victoria Police was the top agency in Victoria by number of OVIC review applications received in 2024-25, ahead of the Departments of Families, Fairness and Housing; Justice and Community Safety; Transport and Planning; and Energy, Environment and Climate Action.

The nine-month average wait is materially worse than the figure that prompted OVIC's own formal investigation. In 2021, OVIC's Commissioner conducted an own-motion investigation into FOI timeliness at five Victorian agencies, including Victoria Police.

That investigation found applicants were waiting an average of six months. Then-Commissioner Sven Bluemmel called it "an unacceptable level of delay" that "deprives Victorians of an important right."

A follow-up report twelve months later found delays had not improved at Victoria Police, the Department of Justice and Community Safety, or the Department of Transport.

The current nine-month figure represents a 50 per cent blow-out beyond the benchmark OVIC had already condemned. No formal enforcement action against Victoria Police has resolved the situation in the intervening period.

Reform stuck in review

Victoria's FOI legislation is also under question at the parliamentary level. The Victorian Parliament's Integrity and Oversight Committee tabled an inquiry report in September 2024 finding that the Freedom of Information Act 1982 was no longer fit for purpose.

The committee called for the Act to be replaced with a modern right-to-information framework. The Victorian Government's subsequent response indicated it needed more time to assess the recommendations, citing complex policy, legal, cultural and resourcing issues.

OVIC said it welcomed the committee's recommendations but would continue to push for improved agency practices in the meantime.

The Right to Know platform, which publishes FOI requests made to Australian public sector agencies, shows Victoria Police responses routinely contain the standard advisory that the office "is experiencing considerable delays in responding to requests and correspondence."

Cyber Resilience Foundations Still Matter in the AI Era: AvePoint

Australian and New Zealand organisations are rethinking cyber resilience and information governance as multi-cloud and multi-SaaS sprawl expands the attack surface and AI agents introduce a new category of risk, AvePoint ANZ Managing Director Max McNamara told the Gartner IT Infrastructure, Operations & Cloud Strategies Conference 2026 in Sydney.

In a session titled "*Cyber Resilience, Governance, and Information Lifecycle Management*", McNamara framed the discussion around four foundational questions: where critical data lives and whether it is protected; whether it could be restored if it became unavailable; whether organisations know where the risk sits across their systems; and what guardrails are in place to prevent and recover from incidents involving AI.

"The protection plane is becoming large," McNamara said. "It's not just Azure, AWS, M365. What about identity? Okta, Entra ID. What about your customer data, Salesforce, CRM? What about ServiceNow? How are you managing workflow data, application data?"

One large Australian aged care provider was the first local example McNamara cited. The provider had been protecting its Microsoft ecosystem, including Azure, M365, Power Platform, identities and Entra ID, through a mix of first- and third-party tools, which McNamara said was time consuming and difficult to maintain.

"They've been able to consolidate that into one view," he said. "They have the ability now to create and maintain those background jobs and that information in one place."

With the backup foundation in place, the provider is now turning its attention to other systems of record that have historically been overlooked. ServiceNow is next on the agenda, with McNamara noting it has evolved from a traditional ITSM platform into a repository for business-critical information. "That wouldn't have even been a conversation if this wasn't already in control."

Across the Tasman, a large New Zealand construction business was cited as an example of the operational gains organisations can extract once they have a unified view of their protection estate.

McNamara said the organisation estimated that insights drawn from AvePoint's "command centre" had saved around 2,000 hours, or roughly 83 days, that had previously been spent manually building reports to understand what was protected, what had exceptions, and what did not.

"A lot of that time was spent being reactive. Now it's proactive," McNamara said. The reclaimed effort has been redirected into resolving exceptions, reviewing scopes and building out containerisation.

Large enterprises cut restore times by 40%

McNamara said two large-scale AvePoint enterprise customers had reduced the time required to identify and restore data by around 40 per cent compared with the legacy platforms they previously relied on.

He stressed that real-world restoration in most enterprises is rarely a full-environment rebuild.

"It's going to be much more granular, a much more surgical toolkit to help them identify the data they're looking for, whether that's a user, a mailbox, a file, and be able to restore that back in place or out of place, depending on the requirement."

McNamara also described an approach to priority-based recovery, which allows organisations to define which groups, teams or user types are restored first in the event of an incident.

"If there was to be an outage or incident, you probably want your executives, your C-suite, to have access to their data and information before maybe someone six or seven levels down in the organisation."

He added that AvePoint had recently embedded Microsoft's first-party Backup for 365 service inside its own platform to provide rapid restore within the native experience for certain workloads.

Organisations shrink the surface area

Shifting the conversation from reactive recovery to proactive information governance, McNamara highlighted examples from a consumer goods company and an Australian healthcare provider as organisations using information lifecycle management to reduce risk before incidents occur. At the consumer goods company, the focus is on archiving redundant content out of production. McNamara said this both reduces Microsoft cloud storage costs and ensures the data that remains in production is relevant rather than obsolete or trivial, shrinking the surface area available to attackers.

The healthcare provider, with more onerous compliance obligations, is rolling out retention and disposition rules so that content no longer required is removed from production and from backups as well.

"That's two examples of organisations moving from reactive cyber resilience to a more proactive position, managing the data in the environment before it gets to an incident, to reduce the surface area and reduce the control plane we're worried about," McNamara said.

Shadow AI: the new resilience challenge

McNamara reserved his strongest concern for what he described as a shadow AI problem emerging as enterprises deploy more agents.

When an agent is spun up, he explained, it typically requires an enterprise application to access Microsoft Graph, email and SharePoint. When the agent is later retired or deleted, the enterprise application often remains in place with its permissions intact, with no clear ownership or monitoring.

"It's still got the permissions. It's still set up for the use case it was set up for. We don't know who's accessing it. We don't know what risk there is. And ultimately, no one's monitoring it. So we don't even know if it's a problem or not."

He argued that AI agents should be governed with the same rigour as human users. Organisations need to know what agents exist, what data those agents can reach, who is accessing them, which of that data is sensitive, and what actions are being taken with it. The same questions, he noted, apply across Google, OpenAI and Anthropic's Claude, not just Microsoft. "It's becoming a challenge. It's not just limited to one vendor."

8 Cyber Priorities for Enterprise in 2026

Artificial intelligence, post-quantum cryptography and the proliferation of non-human digital identities are driving a fundamental shift in enterprise cybersecurity risk, according to a new report by KPMG International.

The *Cybersecurity Considerations 2026* report - drawing on insights from more than 20 KPMG cyber experts and senior leaders from Google, Microsoft, Palo Alto Networks, and ServiceNow - identifies eight priority areas demanding urgent attention from enterprise security leaders.

Non-human identities - including AI agents, service accounts, machine credentials and automated workloads - now outnumber human users by a significant margin. According to research cited in the report, machine identities outnumber humans by more than 80 to 1. Traditional identity and access management (IAM) controls cannot keep pace with the scale and autonomy of these entities.

The report warns that many organisations lack a clear inventory of non-human agents and cannot distinguish between human identities, machine identities, and AI agents. Some agents are being created by other agents, then deleted quickly, leaving little or no trace.

"Tomorrow's identity and access management must evolve beyond static controls into an intelligent, adaptive ecosystem - where realtime decisions and automated policy orchestration appropriately govern and manage the scale of both human and non-human identities," said Juan Manuel Zarzuelo Diaz, Global Digital Identity Leader, KPMG Spain

The report notes that 61 per cent of US companies are not yet comfortable with autonomous agents and will require human-in-the-loop oversight, citing KPMG's *AI Quarterly Pulse Survey* (Q3 2025). A separate finding reported that 59 per cent of companies experienced a data breach caused by a third party in the past 12 months.

The report also highlights a looming workforce challenge. As agentic AI takes on an increasing share of intelligence-driven security tasks - including in the security operations centre (SOC), compliance, and identity management - security teams must acquire new skills to oversee, validate and govern autonomous systems.

In KPMG's *Global Tech Report 2026*, 92 per cent of technology executives say managing AI agents will become an essential skill within the next five years.

"An analogy is the autopilot on a plane, which doesn't replace human pilots, but makes their job more effective and safe. By putting a whole bunch of automation workflows and safeguards and other complex systems around the human delivery, you can drive a much higher degree of efficacy than humans could ever possibly get to on their own," notes Chris Corde, Head of Product, Security Operations, Google

The report calls for organisations to establish cross-departmental governance committees to oversee AI adoption, embed security by design in agent deployment, and implement continuous monitoring including red-teaming exercises.

Safeguarding AI systems is identified as a strategic imperative, not merely a technical challenge. The report warns that over half (54 per cent) of respondents in a global study of nearly 50,000 people across 47 countries - conducted by KPMG and the University of Melbourne - say they are wary about trusting AI.

Post-quantum cryptography (PQC) migration is described as unavoidable, with the transition expected to span multiple years and requiring strategic planning across the entire IT estate. For sectors such as finance and defence, the report characterises the challenge as existential.

The US National Institute of Standards and Technology (NIST) has published PQC standardisation milestones extending into the coming decade. Australian attention was drawn to the issue when Reserve Bank of Australia Governor Michelle Bullock called for proactive collaboration throughout the financial sector in an October 2025 article in the Australian Financial Review.

The report urges organisations to conduct quantum risk assessments, build cryptographic inventories, and update procurement standards to require vendor compliance with quantum-safe requirements.

The full report is available here.

APRA Threatens Action Over AI Governance Failures

APRA has told banks, insurers and superannuation trustees that governance, risk management and operational resilience practices are failing to keep pace with the speed and complexity of AI adoption - and has threatened enforcement action against entities that do not address the gaps.

The Australian Prudential Regulation Authority (APRA) issued the warning in a letter to all regulated entities on 30 April 2026, publishing findings from a targeted supervisory review of large financial institutions conducted in late 2025.

APRA Member Therese McCarthy Hockey said the regulator had observed a fundamental disconnect between AI deployment and governance capability.

"What we've observed from our supervisory engagement is that while AI adoption is continuing apace, the systems and processes required to safely govern its use aren't keeping up. Likewise, the speed at which entities can identify and patch vulnerabilities needs to operate much faster, commensurate with the AI-accelerated threat.

"APRA is also engaging across the sector on the potential for increased cyber threats from high capability AI frontier models such as Anthropic Mythos."

APRA found that while all reviewed entities were actively adopting AI, few had operationalised governance in practice. The letter states that most entities "recognise that existing prudential standards apply to AI risk" but have tended to treat it as "just another technology" - missing key differences such as adaptive model behaviour, bias and data privacy risks specific to predictive systems.

"Many entities are already trialling or introducing AI capabilities in areas such as software engineering, claims triage, loan application processing, fraud and scam disruption, customer interaction and insight generation. However, governance has not matured at the same pace."

Boards were identified as a particular weakness. APRA observed strong interest in AI's commercial upside but said "many Boards are still developing the technical literacy required to provide effective challenge on AI related risks and oversight." APRA also noted an over-reliance on vendor presentations without sufficient examination of model behaviour risks and implications for critical operations.

Cyber Threats Expanding Faster Than Defences

APRA found that AI adoption is materially changing the cyber threat landscape, creating new attack pathways and enabling faster, more coordinated attacks. Specific attack vectors identified by APRA include prompt injection, data leakage, insecure integrations, exploit injection and the manipulation or misuse of autonomous AI agents.

The letter states that AI "can shorten the attack cycle and increase speed, coordination and impact." APRA is separately engaging across the sector on the threat posed by high-capability frontier AI models, noting current Australian Signals Directorate (ASD) advice on frontier models.

APRA identified critical gaps in identity and access management - noting that "capabilities have not yet adjusted to non-human actors such as AI agents." The volume and speed of AI-assisted software development is also straining existing change and release management controls.

Shadow AI use inside financial institutions was flagged as a direct governance concern. APRA observed that entities were relying primarily on "policy direction or detective, after-the-fact measures, rather than enforceable technical restrictions or robust preventative controls" to manage staff use of unapproved AI tools.

Supplier Concentration and Assurance Gaps

APRA found some entities heavily dependent on a single vendor for multiple AI use cases. Few had demonstrated "robust contingency planning or tested exit and substitution strategies for critical AI providers." Contractual arrangements frequently lacked provisions for audit rights, model update notifications and incident reporting.

The letter also highlighted opacity in AI supply chains. When AI capabilities are embedded within broader software platforms or developer tooling, "upstream dependencies such as foundation models, training data sources and fourth party service providers are opaque" - limiting entities' ability to assess model performance, bias, resilience and security independently.

On assurance, APRA found that existing internal audit and risk functions lack the specialist skills and tools required to assess AI systems - particularly where agentic behaviour, automated decision-making or AI-assisted code generation are involved. APRA observed that entities were relying on "point in time and sample-based assurance methods" despite these being "ill-suited to probabilistic models that learn, adapt and degrade over time."

The full APRA letter and executive management attachment are available at apra.gov.au

NSW Treasury Downgrade Leaves Governance Questions Unanswered

The NSW Government has downgraded the "significant cyber incident" declared following an alleged insider data theft at NSW Treasury, with the state's Chief Cyber Security Officer confirming the breach has been contained and remediation measures are in place. The downgrade does not close the matter – criminal proceedings, an internal investigation and continuing legal reviews of potential procurement impacts remain unresolved.

NSW Chief Cyber Security Officer Marie Patane issued the downgrade on 4 May 2026, following confirmation by the government taskforce established to manage the incident. The taskforce confirmed the incident is now in the recovery phase and that agencies had implemented appropriate remediation measures.

A review of potential impacts on active and past government procurements has so far found no project has been adversely affected - though the government qualified this finding, noting that legal reviews are continuing.

A 45-year-old NSW Treasury employee was arrested on 20 April 2026 and charged with accessing and modifying restricted data held in a computer. He is alleged to have transferred more than 5,600 sensitive government documents to an external server between 10 and 14 April 2026.

The documents spanned multiple NSW Government departments and included confidential commercial and financial information relating to current and past government procurement negotiations and private sector transactions.

Under the NSW Government's cyber security framework, a "significant cyber incident" declaration triggers a coordinated whole-of-government response led by the NSW Chief Cyber Security Officer.

The classification in this case was made on the basis of the sensitivity of the data involved and its reach across multiple agencies, rather than the source of the breach - the alleged exfiltration was an insider action rather than an external attack. Cyber Security NSW is required to lead the government's response to significant incidents, coordinating with NSW Police, affected agencies and emergency management bodies. The downgrade does not have a fixed formal definition but signals that the active incident response phase has concluded and standard agency-level recovery processes are sufficient.

The government's downgrade announcement does not disclose what specific remediation measures have been implemented, nor does it address the underlying governance and access control conditions that allowed more than 5,600 documents to be transferred over four days before detection.

NSW Treasurer Daniel Mookhey said at the time of the initial disclosure that the incident required a re-examination of "every system that applies to NSW Treasury." No public update on the outcome of that review has been released.

The procurement integrity finding is the most consequential for the government's commercial relationships. The accused worked in Treasury's commercial team - the unit that oversees major government transactions, procurement and private sector negotiations. Treasurer Mookhey described the stolen files as covering "current government negotiations, previous government negotiations, and interactions" with the private sector.

The government's finding that no project has been "adversely affected" is qualified by "efforts to date" - legal reviews are continuing, and the determination is not described as final. Whether third-party commercial information was among the documents involved has not been confirmed.

The accused was granted conditional bail and is due to appear at Downing Centre Local Court on 3 June 2026. NSW Police have said they believe all the allegedly stolen data has been located and secured, and that there was no external compromise to Treasury's systems. The matter remains before the courts. The government has confirmed a parallel internal investigation is also continuing.

The alleged transfers occurred between 10 and 14 April 2026. NSW Treasury reported the matter to NSW Police on 19 April - nine days after the transfers allegedly began. The government has not publicly stated when internal monitoring detected the breach relative to when the transfers started, nor confirmed the detection mechanism that identified the exfiltration.

Treasury's authorised document repositories include SharePoint, Content Manager, Aurion and TechnologyOne, according to the agency's Information Governance Framework.

The Treasury breach is the latest in a series of insider-related and data exposure incidents involving NSW Government agencies. In October 2025, a government contractor processed a spreadsheet containing personal and health data of up to 3,000 flood victims through ChatGPT.

In September 2025, nearly 600 medical staff had data potentially exposed after NSW Health left confidential documents publicly accessible online.

Earlier in 2025, a man was charged over the alleged access of nearly 9,000 sensitive NSW court documents, including apprehended violence orders and details of minors. The Australian Information Commissioner has noted that insider threats account for a significant proportion of data breach notifications received under Australia's Notifiable Data Breaches scheme, with public sector entities representing a disproportionate share of government-related notifications.

Terabit Attack Exposes Australia's Botnet Blind Spot

A terabit-scale DDoS attack took Australia's largest privately-owned web host offline on Saturday 23 May. The takedown laid bare a regulatory gap that forces enterprise and government customers to fund the country's DDoS defences themselves.

VentraIP and its wholesale sister Synergy Wholesale, both operated by Melbourne-based Nexigen Digital, effectively disappeared from the internet for several hours.

Cheyne Jonstone, Nexigen co-founder and non-executive director, told IDM the Distributed Denial of Service (DDoS) attack likely exceeded the terabit mark. A DDoS attack uses large numbers of compromised devices to flood a target with junk traffic until its internet connections are saturated and it falls offline.

The compromised devices that did the flooding sat in suburban Australian living rooms. A terabit per second is roughly the volume of one hundred thousand simultaneous high-definition Netflix streams, all fired at a single hosting provider.

Eight days earlier, Brisbane-based Binary Lane suffered what it called the largest DDoS attack it had observed. ASD's Australian Cyber Security Centre is engaged with that incident under an open case. Jonstone said Nexigen had been told the source of both attacks was the same.

Asked if those responsible had requested a ransom, Jonstone noted, "We had brief communications with them," but did not provide details.

Residential botnets in the Aisuru and Kimwolf families now exceed one to four million compromised devices globally, almost all of them consumer routers, IP cameras and IoT equipment.

A terabit-scale attack on Australia's largest privately-owned web host exposed a regulatory gap no agency will claim. The bill falls on enterprise and government.

The novelty is the volume now possible. Jonstone pointed to the structural shift. "Traditionally, these attacks came from compromised servers and usually hosted from providers that have outbound mitigation," he told IDM.

"This was completely different, and with NBN services allowing for far greater upload capacity than previous DSL services, the ability to flood our peering and transit networks was far easier."

ADSL2-plus delivered around one megabit upstream. NBN HFC and FTTP routinely deliver 40 to 50. The same botnet, on NBN, is roughly 50 times the weapon it was on copper.

To survive the multi-hour onslaught, Nexigen was forced to emergency-onboard a third upstream provider, GSL Networks, to inject traffic-scrubbing capacity mid-attack. Nexigen has put temporary mitigation in place via a continued arrangement with GSL and is reviewing permanent options.

"It is clear the threat profile has changed," Jonstone said.

The DDoS Tax

No Australian law requires carriers, wholesalers or retail service providers to detect or block outbound malicious traffic from compromised consumer devices. The cost of mitigation falls entirely downstream.

Hosting providers absorb the first hit. They contract premium scrubbing from the likes of Cloudflare, Akamai and Micron21, then pass the cost to enterprise and government tenants through hosting fees. Organisations that self-host or run private or hybrid cloud pay the same vendors directly.

It is, in effect, a DDoS Tax on the customer side of the network, levied to compensate for the absence of any obligation on the carrier side.

Nobody Will Own It

When IDM put the question to the agencies that might own it, none did. The Department of Home Affairs and the Australian Communications and Media Authority refused to respond to inquiries on whether any Australian law, regulation or industry code imposes an affirmative duty on any party in the access-network supply chain to detect, notify about, or act on outbound DDoS traffic from compromised customer premises equipment.

Home Affairs cited SOCI Act incident-reporting obligations that apply to victims of attacks, not to network operators in respect of traffic originating from their own customer base. ACMA directed questions to Home Affairs. The Australian Signals Directorate confirmed on background that identifying regulatory gaps is not its function. The Australian Telecommunications Association advised the questions were outside its jurisdiction.

No body identified any instrument that addresses who, if anyone, in the wholesale-access, retail-access or backbone layer carries an obligation in respect of outbound attack traffic. None indicated the question is under active consideration.

Legally, NBN Co and the major RSPs are acting entirely within the law. The *Telecommunications Act 1997* and the *Security of Critical Infrastructure Act 2018* require carriers to report attacks against their own infrastructure, but neither imposes any duty to inspect, flag or suppress outbound malicious traffic from customer endpoints.

The closest existing instrument is the Australian Internet Security Initiative run by ASD's ACSC, which sends daily reports to participating retail providers identifying compromised IPs. Participation is voluntary, action is voluntary, and the program bypasses NBN Co entirely.

Can NBN Co. Act on Attack Traffic?

A central question raised by the VentraIP attack is whether NBN Co, as the wholesale access operator with a connection to every NBN-served premise, has any technical capability to detect outbound traffic consistent with botnet activity, and whether any obligation requires it to use that capability.

No privacy or interception law appears to prevent outbound anomaly detection. No regulator currently requires it.

NBN Co responded on background, stating that as a Layer 2 wholesale network provider, DDoS attacks at OSI network layers 3, 4 and 7 are outside its visibility, and that responsibility for detection and mitigation lies with the ASN or IP range owners, typically the relevant RSPs or hosting providers. It added that while it constantly monitors network performance, it does not undertake deep packet inspection.

Professor Vijay Sivaraman, Professor of Telecommunications at UNSW Sydney and co-founder of network telemetry firm Canopus Networks, confirmed NBN Co's Layer 2 characterisation is technically correct, but identified the RSP layer as the more practical point of intervention. The Retail Service Providers (RSPs) are the companies such as Telstra, Optus, TPG and iiNet that buy wholesale access from NBN Co and sell internet services directly to residential and business customers.

"NBN is a layer-2 provider, which means its job is to deliver Ethernet frames from the router in the consumer's to the ISP's router located at the point-of-interconnect (POI) exchange. NBN is not meant to be looking at the content of the frames, in fact not even the IP addresses contained in the IP packets within the frames. Therefore, in my view the NBN cannot be expected to provide DDoS scrubbing," Professor Sivaraman told IDM.

"The RSP is probably in a better position to tackle this issue, both in preventing spoofed outbound traffic (which is easy), as well as blocking inbound volumetric attacks (which can be harder as it requires the Netflow/IPFIX extraction and analysis)."

Professor Sivaraman identified two distinct problems at the RSP layer, both technically addressable and both currently unregulated. Source address validation, filtering outbound packets with falsified source IPs, is described by the IETF's BCP38 standard as straightforward to implement. Detection of high-volume traffic from legitimately assigned addresses, as appears to have occurred in the VentraIP attack, requires flow telemetry analysis. No Australian law or instrument currently requires any RSP to implement either measure.

The Installed Base

The *Cyber Security Act 2024* and its Security Standards for Smart Devices Rules took effect on 4 March 2026, requiring new consumer IoT devices in Australia to ship without

(Continued over)

Terabit Attack Exposes Australia's Botnet Blind Spot

(Continued)

default passwords, with a defined support period and a vulnerability disclosure process. Obligations sit with manufacturers and apply only to devices made from 4 March 2026 forward.

But the millions of routers, cameras and smart-home devices already in Australian homes, the exact threat surface that crippled VentraIP, are outside scope and will remain so indefinitely.

Jonstone said the answer has to start with awareness. "I do believe there needs to be far more education to home and small business users about the potential dangers of compromised devices existing on their home or business networks."

"And yes, I believe this will cost end users in the long term as network operators will need to shoulder significant capex and opex to prevent these attacks from happening, especially with the lack of education to end users."

"As a provider of web hosting services, we have tried for nearly two decades to educate our own customers on the importance of keeping their websites up to date and secure, in an environment that we manage. I can only imagine how difficult it would be for an RSP to educate home users as to why connecting everything that comes with wifi to their network," Jonstone said.

An Alternative: Radical Transparency

Lesley Seebeck, an independent consultant and Honorary Professor at the Australian National University, who previously headed the ANU Cyber Institute and has held senior executive roles across the federal government, endorsed more action but urged caution about additional regulatory controls, citing risks around resourcing, capability and surveillance.

"An alternative approach could be one of radical transparency, for example: rather than a closely held feed of classified threat intelligence, ACMA/ACSC could maintain a public, real-time dashboard showing volume of malicious traffic leaving Australian ASNs; focusing on manufacturers, requiring secure-by-design; oversight panels independent of the intelligence community, encouraging (or requiring) acceptance of open-source router firmware and network micro-segmentation for IoT devices," said Seebeck.

An Autonomous System Number (ASN) is a unique identifier assigned to a network or group of networks under a single administrative control, such as an ISP or large enterprise, that exchanges routing information with other networks on the internet. In practice, each major telco and RSP operates one or more ASNs. A public dashboard showing malicious traffic volumes by ASN would effectively name which providers' networks are generating the most outbound attack traffic, creating a reputational incentive to act even in the absence of a legal obligation.

Inside the Corporate Boundary

The threat extends well beyond hosting providers. Australian enterprises and government agencies cannot rely on the wholesale or retail access layer to filter outbound botnet traffic, so the cost of defending against attacks sourced from compromised residential connections falls entirely on the targets.

Hybrid working has extended the problem inside the corporate boundary. An unpatched router on a staff member's home connection sits in the network path between sensitive data and the public internet, outside the Cyber Security Act rules, outside any active cleanup framework, and outside the responsibility of the retail provider that may have supplied it. The same logic covers smart TVs, networked printers and conference-room IoT in the office.

Across the Tasman

The gap is wider still in New Zealand. The country has no consumer device security regime equivalent to Australia's smart device rules, and no statutory duty on its retail providers to do anything about outbound attack traffic.

The *Telecommunications (Interception Capability and Security) Act 2013* covers core network security but says nothing about customer equipment or outbound attack traffic. The 2021-22 DDoS campaign that took Vocus, Kiwibank, ANZ and several public services offline drew substantially on compromised consumer hardware. A compromised router in suburban Auckland can contribute to an attack on an Australian VPN gateway as easily as one in Sydney.

The Means Exist, the Obligation Does Not

The ASD/ACSC Annual Cyber Threat Report 2024-2025 recorded a 280 per cent year-on-year increase in DDoS incidents handled by the ACSC and acknowledged a growing residential botnet attack surface, but does not address carrier-layer obligations. The government's 2023-2030 Australian Cyber Security Strategy does not address residential botnet sourcing or carrier-layer visibility as distinct policy problems.

No parliamentary inquiry, ACMA consultation, Communications Alliance code, or ACSC guidance addresses the question of who in the network supply chain is responsible for outbound attack traffic from Australian residential connections.

The technical means to detect and act on that traffic exist at both the NBN wholesale layer and the RSP layer. The obligation to use them does not.

NZ Health Breach Prompts Privacy Reforms

New Zealand's Privacy Commissioner has found both a major private patient portal vendor and the country's central public health agency in breach of patient privacy laws, following a December 2025 cyber incident that exposed the sensitive records of nearly 100,000 people.

The Phase 1 inquiry report revealed that 99,416 patients had documents stolen from the "My Health Documents" module of the Manage My Health (MMH) portal. While initial estimates feared up to 126,000 people were impacted, the confirmed figure remains a significant breach of sensitive health data.

The incident has triggered calls for New Zealand to overhaul its privacy laws to mirror aspects of Australia's regulatory models.

Manage My Health (MMH is a private, commercial third-party platform that acts as a patient portal and allows patients to view records and manage health data linked to their primary care providers. Health NZ runs New Zealand's national public health system.

NZ Privacy Commissioner Michael Webster announced he will issue formal compliance notices to both organisations under section 123 of New Zealand's *Privacy Act 2020*. Webster described compliance notices as "the strongest tool I currently have available to me to respond to serious privacy breaches".

The bulk of the crisis was concentrated regionally, with around 91% of affected patients residing in Northland. This disproportionate impact was due to a unique local arrangement where Health NZ routinely transmitted hospital records directly to patients via the private MMH portal.

The inquiry concluded the breach was a cascading failure rather than a single technical glitch. Key vulnerabilities identified in the report include:

- **MFA Left Optional:** Multi-factor authentication was available on the MMH platform but was optional for users rather than strictly mandated.
- **Inadequate Access Controls:** Security controls were insufficiently effective, allowing hackers to use a *single* stolen patient account to systematically extract documents belonging to thousands of other patients.
- **Detection Failures:** MMH's data leakage protection was deficient, and its internal systems failed to detect the hackers. The vendor only learned of the breach when alerted by Health NZ.
- **Unaddressed Risks:** Earlier security testing had flagged recurring access control and application security risks, but these themes were "not adequately addressed at the time of the breach".

Health NZ was also heavily criticised for poor vendor governance over what the report termed a "novel and potentially precedent-setting digital project".

The fallout from the breach has prompted the inquiry to recommend structural and legislative changes that draw heavily on Australian and European frameworks:

The report recommends that the NZ Ministry of Health establish a centralised program to independently verify that health tech vendors meet rigorous security standards. The inquiry explicitly cited Australia's My Health Records Act registration system as a successful, comparable model for centralised vendor assurance.

Healthcare Data Fuels Cybercrime Economy

Stolen healthcare data is being traded through a mature global underground economy spanning ransomware groups, access brokers and fraud marketplaces, new TrendAI research reveals.

Over 12 months to February 2026, researchers analysed 7,779 underground forum posts, 21,813 dark web marketplace listings and 95 ransomware leak sites linked to healthcare cybercrime.

Ransomware-related data sales accounted for 36.3% of marketplace activity. Double extortion, where attackers steal data before encrypting systems and threaten to publish it, is now standard practice.

Andrew Philp, Field CISO ANZ at TrendAI, said supply chain compromises involving healthcare software vendors are becoming a major risk multiplier for the sector.

"Patient data is a lucrative target for cybercriminals. Health data is permanent, deeply sensitive and highly reusable, with a single breach creating long-term consequences for individuals, healthcare providers and the wider health ecosystem. The 2024 MediSecure cyber security incident alone saw private data from 12.9 million Australians breached," Philp said.

"This research reinforces why healthcare providers continue to be under the microscope of regulators. Stolen health data is prime currency within the broader underground economy, fuelling criminal activity and creating a ripple-effect across industry and government - with a significant cost for inaction with multi-million dollar fines handed down for healthcare data breaches in recent years."

The research documents a segmented criminal supply chain. Initial access brokers sell network entry points for as little as $US100, feeding ransomware-as-a-service operations.

Pricing follows a clear hierarchy. Small clinic datasets sell for between $US65 and $US400, while bulk medical databases command $US1,000 to $US8,000. Ransom demands against healthcare organisations reach $US500,000.

Fake medical documentation, including fraudulent doctor's notes, disability certifications and sick leave paperwork, sells from $US25.

A key finding for governance and risk teams is the growing targeting of electronic health record (EHR) and electronic medical record (EMR) software vendors. Compromising a single vendor can expose data from hundreds of downstream healthcare practices.

The full report is available here.

FTI Consulting Bolsters Governance Team

Global advisory firm FTI Consulting has responded to rising client demand by appointing a dedicated Australia leader for its Information Governance, Privacy and Security practice.

Kelly Henney has joined FTI Consulting as Senior Managing Director and Australia Leader within the firm's Technology segment, based in Sydney. She specialises in AI governance, data privacy, digital risk, information security and data resilience.

Henney brings more than a decade of experience in digital risk, including data breach response, regulatory enforcement, investigations and enterprise-wide governance programs. She previously led privacy and data protection advisory at a Big Four consulting firm and has practised as a corporate lawyer in banking and finance.

"AI is moving so quickly that many organisations are struggling with how to develop and implement adaptable, actionable governance frameworks," Henney said. "The keys to solving this challenge are ongoing education about AI risk and capability, alongside rigorous testing and oversight processes that are built in at the entry point and throughout deployment."

Veeral Gosalia, Senior Managing Director and Leader of FTI Technology in Asia Pacific, said Australia's regulatory environment is becoming more demanding. "Australia's privacy and AI regulatory landscape is continually evolving to keep pace with the digital landscape, making compliance an ongoing challenge for our clients," Gosalia said.

The appointment comes as enterprises confront a series of rolling compliance deadlines. A statutory tort for serious invasions of privacy took effect on 10 June 2025, giving individuals direct legal recourse against organisations that intentionally or recklessly misuse personal information.

From 10 December 2026, all APP entities must disclose in their privacy policies how personal information is used in automated decision-making - a requirement with direct implications for organisations deploying AI systems. Non-compliance carries civil penalties of up to $50 million, three times the benefit obtained, or 30 per cent of adjusted turnover, whichever is greater.

FTI Consulting has published an overview of the reforms and their compliance implications at https://www.fticonsulting.com/insights/articles/australian-privacy-law-reforms-take-effect.

Kiwi Firms to Disclose AI Use In Australia

New Zealand law firm Buddle Findlay has warned that Kiwi businesses that hold or handle the personal information of individuals in Australia must prepare for strict new transparency rules regarding automated decision-making.

From 10 December 2026, amendments to the Australian Privacy Act will force organisations to reveal when they use computer programs to make significant decisions about individuals. These changes target systems that rely on personal information to screen job applicants, assess credit scores, or detect fraud.

For New Zealand's CIOs and GRC managers, the update highlights a growing regulatory gap. While Australia is moving toward mandatory AI disclosure, New Zealand's Privacy Act 2020 remains silent on the issue. This means firms relying solely on local compliance will likely fall short of Australian legal standards.

According to Buddle Findlay, the new rules have a long reach. Any New Zealand entity that does business in Australia or handles the data of people living there may be classified as an "APP entity".

Under sections APP 1.7 to 1.9, these organisations must clearly describe in their privacy policies:

- The kinds of personal information used in the operation of such programs

- The kinds of decisions made solely by those programs

- The kinds of decisions for which those programs perform a step substantially and directly related to making the decision.

"For New Zealand organisations operating in or across the Australian market, the practical consequence is clear: compliance with APP 1.7 in Australia is a standalone obligation that must be addressed on its own terms. It cannot currently be satisfied by, compliance with New Zealand's Privacy Act. "

The law firm has published 4 practical steps that should be followed by in-house legal teams and privacy officers at New Zealand organisations with Australian operations or exposure.

Glasswing flies in to ANZ Region

Anthropic has expanded Project Glasswing, its AI-powered cybersecurity vulnerability program, to approximately 150 new organisations across more than 15 countries - with Australia and New Zealand among the nations included in the second cohort, according to reporting by the Financial Times. Anthropic has not publicly confirmed the country list.

The Financial Times reported that the expansion includes organisations from Australia, New Zealand, Canada, France, Germany, Italy, Japan, South Korea, the Netherlands, Spain, Belgium, Sweden, and Switzerland, along with India. Anthropic declined to confirm the countries involved, stating: "We're not disclosing the specific organizations or the list of countries publicly."

Project Glasswing gives participating organisations access to Claude Mythos Preview - Anthropic's most advanced AI model, which is not available to the general public - for the purpose of scanning their codebases for software vulnerabilities before they can be exploited.

The program was launched in early April with around 50 initial partners, including Apple, Microsoft, Nvidia, CrowdStrike, Palo Alto Networks, and the US government. Those early participants identified more than 10,000 high- or critical-severity security vulnerabilities within the first weeks of using the model.

The expanded second cohort broadens the program into sectors not well represented in the initial group, including power, water, healthcare, communications, and hardware. Anthropic said the new partners are predominantly vendors whose codebases underpin systems used by many organisations worldwide, including governments.

The Glasswing announcement coincided with US President Donald Trump signing an executive order titled "Promoting Advanced Artificial Intelligence Innovation and Security" on June 2. The order establishes a voluntary framework under which AI developers are asked to provide the US government with early access to frontier models for up to 30 days prior to broader release, to allow security review.

The order explicitly states it does not create a mandatory licensing, permitting, or approval process for AI models - a position that represents a retreat from an earlier proposal that drew significant industry opposition.

Anthropic framed the Glasswing expansion as an urgent response to the broader trajectory of AI development. The company warned that within six to 12 months, other AI developers are likely to release models with comparable cyber capabilities, potentially without adequate safeguards.

"Cheap, fast AI models with powerful cyber capabilities are around the corner," Anthropic said. "We want Project Glasswing to spur institutions toward operating norms that reflect this reality."

Knosys Secures $A3.8m ANZ Extension

ANZ has signed a $A3.8 million two-year contract extension with Melbourne SaaS provider Knosys for its enterprise knowledge management platform.

The agreement includes an optional third year and requires ANZ to pay the full multi-million-dollar contract value upfront to Knosys.

The extension ensures the continued deployment of KnowledgeIQ, which acts as a centralised and authoritative data management platform within the bank.

The platform functions as a critical System of Record in the artificial intelligence value chain to support reliable organizational decision-making.

It provides a trusted primary repository for specific business data, ensuring data accuracy, integrity, and consistency across internal operational processes.

This renewal follows an interim one-year extension signed in 2025, where both organisations began collaborating on cloud migration and automated features.

"This contract extension reflects our long-standing relationship with ANZ and our ability to support their internal strategic initiatives such as the integration of a new AI assistant into their knowledge management portal," said John Thompson, Knosys Managing Director.

Australia Post taps global AI expertise

Australia Post is partnering with Alpha Level, a next generation AI security company, co-founded by international machine learning expert Dr. Josh Neil, to sharpen its cyber defences and speed up threat detection across its network.

Alpha Level will apply cutting-edge machine learning to help Australia Post efficiently collect, process and analyse billions of data points - boosting both the speed and accuracy of threat identification.

Machine learning - a branch of artificial intelligence - enables systems to rapidly analyse vast data sets, learning patterns and improving decision-making over time without needing explicit programming.

Australia Post Chief Information Security Officer, Adam Cartwright said the partnership with Alpha Level will enhance the effectiveness and efficiency of Australia Post's cybersecurity measures.

"Cyber safety remains a critical focus for Australia Post, and as threats continue to evolve, we're focused on strengthening how we identify and respond to them.

"Each month, our systems generate around four billion data points - from network traffic to security logs - which our team assesses for signs of malicious activity. It's a complex, time-consuming task.

"Machine learning helps by building models that understand what 'good' looks like in that data, allowing us to detect threats faster and more accurately. This not only speeds up detection but strengthens our overall security in an increasingly complex digital environment.

"By bolstering our cyber posture, we're not only protecting Australia Post's infrastructure - we're also helping to lift the security of the thousands of licensed post offices and small businesses that depend on our network to serve their communities," said Mr. Cartwright

With a PhD from Los Alamos National Laboratories, Dr. Neil is a former Principal Data Scientist at Microsoft and specialises in machine learning, statistical mathematics and AI.

Alpha Level Chief Technology Officer, Dr Josh Neil, said the partnership demonstrates how advanced machine learning can strengthen cyber defence at scale.

"We're proud to partner with Australia Post, an organisation that's leading by example in proactive cyber defence.

"By embedding world-class machine learning into Australia Post's security operations, we're helping to elevate threat detection speed and precision - ensuring resilience against the increasingly complex challenges in today's digital landscape," said Dr Neil.

Why the Sovereign Cloud Aspiration is Harder than it Sounds

Gartner analyst Douglas Toombs warns that 80 years of globally interconnected trade will not unbundle neatly in 18 months, and that much of what passes for 'sovereign cloud' today is little more than 'sovereign washing'.

Four American hyperscalers - Amazon Web Services, Microsoft Azure, Google Cloud and Oracle Cloud - between them control somewhere north of 80 per cent of the global enterprise cloud market.

That single fact, more than any political event or piece of legislation, frames every serious conversation about sovereign cloud taking place inside enterprises and government agencies today.

Speaking at the Gartner IT Infrastructure, Operations & Cloud Strategies Conference 2026, Gartner analyst Douglas Toombs laid out an unsparing assessment of just how difficult it will be for any organisation, or indeed any country, to extract itself from that concentration.

Toombs pointed to Boston Consulting Group's long-standing 'rule of three and four', which holds that a stable competitive market never has more than three significant competitors, the largest of which has no more than four times the share of the smallest.

The 2024 enterprise cloud infrastructure and platform services numbers fit the pattern almost too neatly: AWS at roughly US$100 billion, Microsoft at $66 billion and Google at $29 billion, with Alibaba, Huawei, Tencent, Oracle and IBM scrapping over the remainder.

"With 80, 90-plus per cent of market share being just a handful of companies, I don't care what political party you like or dislike or whatever you think, it's just hard from a market dynamics perspective to move that type of thing," said Toombs.

Gartner's own published position, as captured in research note *Critical Insights: Impact of US Federal Policy Changes on International Cloud-Centric Tech Services*, is even blunter: "There are (currently) no suitable non-US alternatives to Amazon Web Services, Microsoft Azure, Google Cloud and Oracle Cloud."

A surge in interest

Despite that uncomfortable arithmetic, the rate at which Gartner clients are now asking about sovereign cloud has exploded. Toombs said inquiry-call volumes across his analyst peers tell a clear story.

"From the end of 2024 through the first quarter of 2025 and then through first quarter of this past year, the growth in interest in sovereign cloud globally was about five-fold. So a significant global increase.

A lot of questions, a lot of conversations. Would we have had that same five-fold increase if the November 2024 US election went slightly differently? Who knows? But we are where we are."

The interest is heaviest in Western Europe but is visible globally. The trigger that comes up in client conversations again and again, Toombs said, is the 2018 US CLOUD Act and the prospect of American law enforcement reaching into data held by American providers regardless of where in the world it sits.

The Act that never really was debated

Toombs offered a brief and slightly disquieting history lesson on the CLOUD Act (Clarifying Lawful Overseas Use of Data) itself. The legislation was originally drafted as a standalone bill, S.2383, sitting on a desk in February 2018. Its text was then lifted wholesale and pasted into the must-pass H.R. 1625, the Consolidated Appropriations Act of 2018, which was signed into law on 23 March 2018.

"The Cloud Act was never debated as one would expect in a legislative body," Toombs told delegates. "Normally we send it to a committee, people argue, make amendments, but in this case a bunch of things written down, sat for a bit, copy and paste, law."

The Act amended 18 US Code §2703, allowing a provider to file a motion to quash a US legal demand only where the subscriber is not a US person and where the disclosure would create a material risk of violating the laws of a 'qualifying foreign government', defined as a country with which the United States has signed an executive agreement under section 2523.

As of March 2026, the US Department of Justice lists only two such agreements: with the United Kingdom and with Australia. "So that's actually good," Toombs noted. "That's less good when I'm on continental Europe giving this presentation."

It is worth noting that Australia signed its CLOUD Act executive agreement in December 2021, while New Zealand has no equivalent. There is no public data on how frequently the Australian agreement has been invoked.

A long graveyard of sovereign ambitions

Sovereign cloud is hardly a new idea. Toombs walked through almost two decades of attempts, most of them now defunct or moribund.

France launched Project Andromède in 2009 with 285 million euros of state backing, producing the Cloudwatt offering.

"You can't log in to Cloudwatt today and provision a virtual machine," Toombs observed.

France tried again in 2016 with Numergy and SFR, again with hundreds of millions of euros behind it. That too has gone.

Microsoft built Microsoft Cloud Germany in 2016 under a data-trustee arrangement with Deutsche Telekom's T-Systems. Two years later, buried in the closing lines of a corporate blog post, Microsoft quietly stopped taking new customers.

The Gaia-X initiative announced in 2020 has, by Toombs' reckoning, produced rather more white papers than running infrastructure.

The IPCEI Next Generation Cloud Infrastructure and Services project approved in December 2023, and the Dutch-led ECOFED federation that followed, remain works in progress.

And it is not only smaller players or European national champions that have come to grief. VMware vCloud Air was eventually offloaded to OVH. HP Helion was shut down less than two years after launch.

Telstra's much-promoted A$800 million cloud play, announced by then-chief executive David Thodey back in 2011, no longer occupies a prominent place in the carrier's portfolio.

"It's not just small companies or non-US companies or whatever. This is a hard market to break into, at least at the cloud infrastructure and platform layers," said Toombs.

Sovereign washing

Toombs reserved some of his sharpest commentary for what he called 'sovereign washing', the practice of dressing up offerings that are still ultimately tethered to US corporate parents in patriotic local livery.

Telstra's 2024 announcement of a sovereign secure cloud built in partnership with AWS is one example.

"There's some in Europe that have been launching there too," he said. "Sure, ultimately it's compute, storage, networking, CPU cycles for the providers. And if that's fine for your customers and your regions that have concerns, great. But it's interesting having watched this for a long time, seeing how the pendulum swings back and forth."

Microsoft's June 2025 announcement of Microsoft 365 Local prompted a similar reaction. The offering, marketed as part of Microsoft's comprehensive sovereign solutions for European organisations, allows customers to deploy productivity workloads like Exchange Server and SharePoint Server in their own data centres.

(Continued over)

"I'm an Exchange administrator from the dot-com era," Toombs said. "I was doing that 25 years ago. How is that anything new?"

The point is that the customers driving the sovereign cloud conversation moved to Microsoft 365 software-as-a-service, with Teams at the centre, during and after the pandemic. They do not want to go back to running Exchange and SharePoint on a Windows server in a cupboard.

The bigger problem: technology sovereignty

Even if an organisation accepts that data residency and operational autonomy can be achieved through a local partner offering, the deeper question of technology sovereignty remains.

Toombs put up a slide breaking the technology stack into eight layers, from hyperscalers down through database, operating system, server virtualisation, servers, network and storage, other hardware and 'other miscellaneous' software vendors.

The US flag dominates almost every row.

"The challenge on the technology sovereignty side is that when you start to think about the entire landscape of vendors you would typically work with to power your applications, a lot of them fall under the US flag. So, if you want to get rid of all of those, that gives you a much smaller sourcing and management landscape to work with."

Gartner's position, drawn from its research note *Digital Sovereignty Is Needed For Future Technological Resilience and Business Outcome*, is unequivocal: "Outside of the US and China, true technological sovereignty cannot be achieved today."

Toombs cited the now-notorious July 2025 exchange in the French Senate, when Microsoft's French general manager and head of legal affairs Anton Carniaux was asked under oath whether the company could guarantee that data on French citizens would not be transmitted to the American government without the explicit agreement of the French government.

His answer, as Toombs paraphrased it, was a single word: "No."

Five strategies, and a warning about getting cornered

Having pulled apart the comforting fictions, Toombs set out five practical strategies organisations can choose from, each with its own profile of risk, complexity and cost.

Shelter-in-place is, as he frankly described it, "do nothing and wait for things to go back to normal". It carries the highest residual risk but the lowest cost, and Toombs said it is in fact where most of his clients currently sit.

Distributed cloud uses hyperscaler-provided hardware (AWS Outposts, Azure Local, Oracle Dedicated Region, Google Distributed Cloud) in customer or local-provider locations. It addresses jurisdictional residency but most variants remain tethered to the mothership for operational and update purposes, and the service catalogues are typically a tiny subset of what is available in the public regions.

Sovereign cloud comes in three flavours: hyperscaler-operated sovereign offerings such as the AWS European Sovereign Cloud launched in January 2026 and the Oracle EU Sovereign Cloud in Frankfurt and Madrid, hyperscaler-supported partner offerings such as Bleu (Capgemini and Orange on top of Microsoft technology) and S3NS (Thales on Google Cloud), and pure local non-hyperscaler offerings.

Toombs cautioned that even the hyperscaler-operated sovereign clouds may remain legally owned by US parent companies, and that local non-hyperscaler offerings frequently lack feature parity.

"You bought the Corvette, you needed a golf cart, but you might have a local provider that's trying to sell it as the skateboard," he said.

Hide-in-plain-sight wraps existing public-cloud deployments in customer-managed encryption keys and confidential computing technologies.

It addresses the data-disclosure risk by ensuring the provider cannot hand over usable data, but does not protect against kill-switch scenarios or service disconnection.

Private, hybrid and multi-cloud spreads workloads across locations to balance risk against innovation, at the cost of significant operational complexity and capital investment, and without solving the problem that some innovations simply do not exist outside the big public clouds.

Underpinning all five is what Toombs described as the single most important conversation any CIO needs to have with their leadership: nailing down the specific trigger events that would prompt an exit, and the timeframe required to execute on it.

"The faster you need to exit, the more you need to spend upfront. Do not let yourself get cornered into the position where leadership tells you the exit needs to happen immediately, but you have not been given the budget or the time to prepare for it," said Toombs.

Geopatriation, not repatriation

Toombs and his colleagues have coined the term 'geopatriation' to distinguish jurisdictionally driven cloud exits from the broader repatriation narrative that has bubbled along for several years.

Gartner, he said, sees little evidence of mass repatriation despite the headlines, and so far only modest signs of actual geopatriation: some forward projects being paused, but very few existing workloads being pulled back.

Part of the reason is the unavoidable competing priority of business continuity. Toombs cited recent incidents in which cloud provider regions in Gulf states were directly impacted by missile strikes.

"We said, sovereignty-wise, it has to stay here, but this data centre is on fire. What do we do?" he said. "It becomes a pick-your-disaster-scenario exercise, and it is going to cost a gazillion dollars."

Toombs closed with five points he urged delegates to take back to their organisations. Consolidation of more than 80 per cent of market share into three companies makes any market exceptionally hard for new entrants to break into. The sovereign cloud landscape is replete with lofty pro

mises and dead offerings. Do not let anyone hand-wave what your exit triggers should be. Demand the appropriate structure, budgets and time, to do it right. And implement whichever of the five key strategies protects your organisation best, at the right investment level.

Or, as Toombs put it borrowing from Thomas Edison: vision without execution is hallucination.

National Framework Targets Fragmented Health Data Standards

The Australian Digital Health Agency (ADHA) concedes adoption of standards remains "sporadic" more than two decades after foundational digital health work began in Australia but claims a new framework will drive uptake.Launching a new *National Framework for Digital Health Standards*, Agency Chief Executive Officer Amanda Cattermole PSM said, "Different organisations have developed and applied standards in isolation, with limited coordination to fit those pieces together across the system."

"The *National Framework for Digital Health Standards* provides clear, practical guidance and aligns governance, standards development and implementation across government agencies, jurisdictions, health services, partners and industry."

The 38-page frameworkc proposes a national governance model, formal standards selection criteria and a new Standards Academy. It targets inconsistent uptake by software vendors, hospitals and health services.

The framework points to SNOMED CT-AU as an example of slow adoption. The former Australian Health Ministers Advisory Committee endorsed it as the national clinical terminology in 2005.

The report states its "implementation and use in clinical information systems and digital health solutions is not nationally consistent".

HL7 v2 messaging has achieved near-universal use in some sectors. Transition to the newer Fast Healthcare Interoperability Resources (FHIR) standard is still in progress across Commonwealth Government systems.

The framework identifies six structural barriers to national adoption.

These include vendor discretion over which standards to support, commercial incentives to build proprietary alternatives and the cost of upgrading legacy clinical systems.

It also blames "limited policy drivers (legislative, financial) to encourage sector-wide implementation of interoperable solutions and standards".

The framework cites the absence of a national governance system for endorsing and adopting information-sharing standards as a key gap.

The Agency says privacy and cyber security risks have constrained the pace at which legacy systems can be modernised.

Business cases must also account for change management, training and software replacement costs.

Five recommendations

The framework sets out five recommendations to be pursued with the Department of Health, Disability and Ageing (DHDA), Standards Development Organisations and jurisdictions:

- Establish an enduring national governance model for standards adoption.
- Strengthen the evidence base for standards benefits.
- Build workforce capability through training and education.
- Enhance communication, collaboration and engagement.
- Provide practical implementation support including catalogues, procurement guidelines and test environments.

A National Roadmap for Digital Health Standards will be developed alongside the framework to signal transition timeframes for sector planning. ADHA points enterprise architects, information managers and procurement teams to several existing tools.

These include the Digital Health Standards Catalogue, Procurement Guidelines, the National Clinical Terminology Service and the Australian Register of Conformity.

The agency has opened a Standards Academy offering free training on SNOMED CT, GS1 and FHIR for clinicians, developers, policymakers and industry. https://www.digitalhealth.gov.au/digital-health-standards/standards-academy

Aged Care Clinical Information System (ACCIS) Standards have also been developed to address recommendations from the Royal Commission into Aged Care Safety and Quality.

The framework includes a conformance and assurance pathway covering clinical safety, cyber security, architecture and design, legal and policy requirements for vendors connecting to national infrastructure.

Mandatory sharing approaches

Agency Chief Digital Officer Peter O'Halloran said the framework is needed as legislated information-sharing obligations take effect.

"As mandatory sharing expands to more key health information, scalable and conformant national digital infrastructure becomes critical to safe, reliable information sharing across the health system," Mr O'Halloran said.

Australia is rolling out the Health Legislation Amendment (Modernising My Health Record - Sharing by Default) Bill 2024. The legislation will require pathology and diagnostic imaging providers to share test results to My Health Record by default.

Mr O'Halloran said pathology views have grown 112 per cent over the past year, from 54 million to 114 million. Diagnostic imaging views are up 72 per cent.

The framework also positions standardised terminology as a control for emerging artificial intelligence risks. It states nationally adopted coding maps reduce "the risk of incorrect or hallucination impacted creation of structured data".

The framework cites overseas examples where legislation has been used to force standards adoption. These include Canada's Bill C-72, which proposes prohibiting data blocking.

The US 21st Century Cures Act mandates FHIR-based certified health IT. NHS England's Data Saves Lives strategy links funding to interoperability conformance.

A 2024 Productivity Commission research paper on leveraging digital technology in healthcare found Australia's health information environment "remains fragmented".

It said modernisation of national digital health infrastructure is critical.

Download the report here.

Funding Squeeze hits OAIC as Privacy Reforms Land

With three major regulatory expansions landing on its desk in the next 13 months, the Office of the Australian Information Commissioner (OAIC) has suffered a funding cut in the 2026 federal budget,

The OAIC has been allocated $36.576 million in the 2026-27 Federal Budget. The figure is down from $39.753 million in 2025-26, representing an 8 per cent cut to the agency's departmental appropriation.

Average staffing levels at the OAIC remain almost flat. The agency is funded for 175 staff in 2026-27, down from 176 in 2025-26.

Privacy Commissioner Carly Kind has previously been candid about OAIC resource constraints. Asked at a public event in 2024 what she would want if she could have anything, Kind replied "an unlimited enforcement budget".

The OAIC is responsible for enforcing three significant Privacy Act expansions arriving between July and December 2026.

From July 1, the Anti-Money Laundering and Counter-Terrorism Financing regime expands to cover new sectors. Lawyers, conveyancers, accountants, real estate professionals and dealers in high-value goods become reporting entities.

Small business reporting entities lose their Privacy Act exemption when they enter the AML/CTF regime. The OAIC estimates this change will affect more than 100,000 small businesses.

The OAIC role expands accordingly. The regulator becomes responsible for privacy oversight of a much larger pool of small business entities.

From 10 December 2026, the Privacy Act will require businesses and government agencies to tell people when they use computer programs to make important decisions about them.

Organisations must update their privacy policies to explain three things. They must list the personal information their automated systems use. They must list the types of decisions those systems make.

Privacy policies must also identify which decisions could significantly affect a person's rights or interests.

The rules cover any computer program that makes a decision, or substantially helps make one. This includes rule-based software, machine learning models and generative AI tools.

Who It Affects

The rules apply to any business or agency already covered by the Privacy Act that uses automated tools to make decisions about people. A decision is in scope if it could reasonably be expected to significantly affect a person's rights or interests.

Examples include a bank using an algorithm to approve or refuse a loan. An insurer using software to set a premium or reject a claim.

A government agency using a system to grant or refuse a benefit is also in scope. So is an employer using AI to screen job applicants, or a healthcare provider using software to support diagnostic decisions.

The rules apply to any decision made on or after 10 December 2026. It does not matter when the algorithm was built or when the data was collected.

The OAIC will enforces the new rules. The Office will publish guidance on what organisations must include in their privacy policies.

A Children's Online Privacy Code is also due to be registered by 10 December 2026.

The OAIC role includes drafting the Code, conducting a 60-day public consultation, and registering and enforcing it. The Office began consultation work in 2025.

NSW Police Force Needs $A493m to Finish Technology Reset

A report by the NSW Auditor-General has found the NSW Police Force needs four extra years to complete its core technology upgrade, requiring $A493 million in additional funding. The agency has spent $155 million while delivering only one of five core systems.

The report covers planning, procurement and governance activities at the NSW Police Force from 2018 to December 2025. The audit concluded the NSW Police Force "has not efficiently or effectively planned and sourced key components to upgrade core policing technology systems."

The program's original 2027 completion date has been pushed to June 2031. The NSW Police Force estimates it needs $78 million more in capital funding and $415 million in recurrent funding to finish the work. By December 2025, the agency had spent 47 per cent of its $328 million capital allocation. Only the forensics and exhibits system has been replaced in full.

The Computerised Operational Policing System (COPS), implemented in 1994, remains in use on outdated mainframe technology using an obsolete programming language. COPS sits at the centre of more than 300 systems and subsystems, with 207 interfaces to internal and external platforms. The audit followed a 20 November 2024 referral from the Law Enforcement Conduct Commission to the Audit Office of New South Wales. The referral concerned the NSW Police Force's administration of the Integrated Policing Operating System (IPOS) project.

Failed contract with overseas start-up

In 2020, the NSW Police Force awarded the prime contract for IPOS to an overseas start-up technology firm. The audit said the supplier scored higher on operational criteria than two competing bidders.

External advisers and gateway reviews had warned decision makers about delivery optimism, resourcing constraints and supplier capability risks. Concerns also covered the firm's financial position and limited experience delivering systems at comparable scale.

In June 2022, the supplier advised it could not deliver a forensics and exhibits system. Minimum viable computer aided dispatch capability would also slip by several years. The NSW Police Force terminated the contract.

Total lost investment reached close to $20 million. The supplier began legal action against the NSW Police Force, with a confidential settlement reached in August 2023. Following contract termination, the NSW Police Force did not maintain effective governance, capability or appropriate financial controls, according to the audit. This slowed decision making and prolonged reliance on legacy systems.

Between 2022 and 2024, the steering committee did not provide consistent business leadership or effective oversight of timelines, budgets or risks.

A consultants' project management review in 2024 found projects were "being managed in silos rather than through an interconnected portfolio." The review also identified "poor project/program financial controls and inconsistency in financial management."

In January 2023, the IPOS program had 191 approved positions, with 82 (43 per cent) vacant. Recruiting programmers with legacy system skills proved especially difficult. The cost of temporary developers rose from $1,000 per day in December 2021 to $1,500 per day in December 2024, equivalent to $345,000 per year.

Program reset and multi-vendor pivot

The NSW Police Force initiated a comprehensive reset in mid-2024, rebranding the work as the Police Technology Program (PTP). It introduced single accountable sponsorship, a phased delivery model, strengthened financial oversight and reinforced independent assurance.

The agency has moved away from a single-vendor approach. Procurement processes in 2024-25 for next-generation CAD and the COPS replacement required tenderers to meet minimum staff and revenue thresholds, effectively excluding start-up firms.

Proposals are accepted only from companies with experience in the Westminster system of government that comply with the Australian Protective Security Policy Framework. A Department of Customer Service health check in November 2025 found the program had "pivoted to a more flexible and scalable delivery model." It rated overall delivery confidence as medium.

Operational cost of legacy systems

The audit detailed operational impacts of running 1994 technology. Officers must enter the same information multiple times and return to stations to input data. Staff often use paper-based processes to overcome problems with COPS.

A 2019 case study highlighted the impact during a $4 million family day care fraud investigation. Some 700 exhibits were manually entered into four different systems. Search warrant documents were entered into five different systems retrospectively, as there was no scope for realtime input.

Updating COPS for a 2023 bushfire reporting change required 110 days of developer time due to legacy system complexity. The mainframe costs approximately $13 million per year in hardware and software. An additional $8 million per year covers software licences.

The 2020 business case assumed mainframe decommissioning would save around $26 million per year from July 2027 onwards. Those savings are no longer expected.

NSW Police Commissioner Mal Lanyon accepted both audit recommendations. In his formal response dated 24 April 2026, he wrote that the reset was designed to limit further exposure and preserve service continuity.

Lanyon said the reforms "represent more than incremental improvement; they constitute a fundamental restructuring of how technology transformation is governed and delivered within the NSW Police Force."

The full NSW Auditor-General's report, Upgrades to core policing technology, is available at audit.nsw.gov.au.

Resetting the Data Governance Narrative

A Joint Statement from John Ladley & Robert S. Seine

For years, we have watched data governance be misrepresented, over-complicated, and – let's be honest – over-cooked. Organizations struggle with ineffective implementations, leadership resists investing in governance, and practitioners find themselves caught between rigid compliance mandates and the realities of how data is used. The problem isn't data governance itself – it's how it's being understood, communicated and applied.

As two long-time advocates for practical, results-driven data governance, we are using our voices, experience, and influence to call for a reset. Our message is not intended to promote our practice, or further interests in a product or service. Candidly, we want the reader to start to think for *themselves* instead of relying on a talking head or vendor-driven scenarios.

The world has changed. AI, automation, real-time data, and transitions in the workforce are forcing organizations to rethink how they manage and use information. "Data driven" appears in annual reports. Yet, we still see governance framed as an IT function, treated like an administrative roadblock, or a set of restrictive policies rather than what it truly is – a business capability that is a catalyst for efficiency, trust, and innovation.

Lessons We Must Acknowledge to Move Forward

Data governance is not a failure (yet) – how we talk about it is. Organizations already handle data informally. These efforts are inconsistent, inefficient, and ineffective. Never forget that governance already exists in your organization so the focus must be on practical refinement of what exists or practical iterations. "Data" governance is an expansion of a long-accepted concept.

Data governance is mis-interpreted – Data governance is widely misunderstood – and often dismissed as stale – because people forget it's not new. At its core, governance simply means the execution and enforcement of authority – it's the check in the checks and balances that ensure your data works. That's why data governance must move out of IT. Keeping it there creates a conflict of interest and undermines fundamental business principles. Governance belongs with the business, where accountability lives.

Failing to convince leadership that data deserves governance is one of our biggest setbacks. We must normalize the idea that formal, monitored behavior around data is just standard business practice – no different than how finances or operations are governed. Governance is already in the job description of CEOs and boards of directors. Without it, an organization's actions can't be verified or trusted. When someone claims, "data governance slows us down," what they're really saying is, "we're okay with producing unmeasurable, un-verifiable results." That mindset is the real threat to progress.

Governance must enable, not obstruct. Governance is often implemented in ways that slow progress rather than accelerate it. Governance must be positioned as an *EFFICIENCY driver* that supports and enables value from all uses of data assets. Technology can help administer and enable governance, but, like providers of other software tools, vendors want to sell their software and will assert themselves at any opportunity.

Leadership will benefit significantly from a business-driven message. The message must be one that is shielded from technical jargon. If we continue to pitch governance in terms of data quality, metadata, lineage, and data catalogs, we will continue to lose the attention of executives. These items are part of a data professional's job. Governance must be framed in terms of revenue growth, risk mitigation, operational efficiency, and AI and regulatory readiness.

Leadership will benefit significantly from the desire to learn new things. There are aspects of Data Governance that echo the sounds of Total Quality Management (TQM). TQM required leadership to learn new methods of operation. Industry data people have shied away from engaging the highest level. We think that needs to change. We must engage. They must also engage.

Change management is a discipline that our industry largely ignores. Data governance requires change. The industry needs to focus on behavior changes and embedding governance into everyday work without unnecessary disruption.

Governance that doesn't drive measurable business value is doomed to fail. Every governance initiative must be directly tied to outcomes like improved customer experience or reduced risk – or it won't stick. This isn't revolutionary; Marketing, HR, and Accounting are all enabling functions held accountable for results. Governance must be no different. Success demands alignment and intentionality, not avoidance.

Using Our Influence to Drive Change

We are done watching data governance be misunderstood and misapplied. So are many of our peers. The call to action must shift the conversation, cut through the noise, and start treating governance like the required business capability it truly is.

As recognized leaders in the industry, we are committed to using our platforms, experience, and networks to reframe how governance is discussed and implemented. Our operating models for DG must facilitate and accelerate organizations' cultures, and readiness. Data governance must be aligned with organization strategy. AI, data-driven decision-making, and digital transformation demand governance that is practical, adaptable, and directly tied to business success.

The call to action is this. Data leaders, executives, vendors, and practitioners need to hear a new narrative. Leadership is the audience. They do not only want to hear about better reports, the need for metadata, or the wonder of a data catalog.

We need to look up from our shoes, look them in the eye, communicate more effectively, and educate. This will mean repeated statements to an unwilling audience. "The world has changed. Leadership needs to monitor data assets and support better data behaviors – or you will never be fully data-driven."

There will be many opportunities in the coming months to weigh in. It is time to stop treating governance as an afterthought, or as an obstacle, a means to make a quick sale, or even as a new, miraculous panacea – and start making it the business-critical, required capability it was always meant to be.

Originally published here.

Robert Seiner is Author of Non-Invasive Data Governance Book Series, The Data Catalyst (Cubed), President & Principal of KIK Consulting & Educational Services, and Adjunct Faculty at Carnegie Mellon University. John Ladkey is Author of Data Governance: How to Design Deploy and Sustain an Effective Data Governance Program.

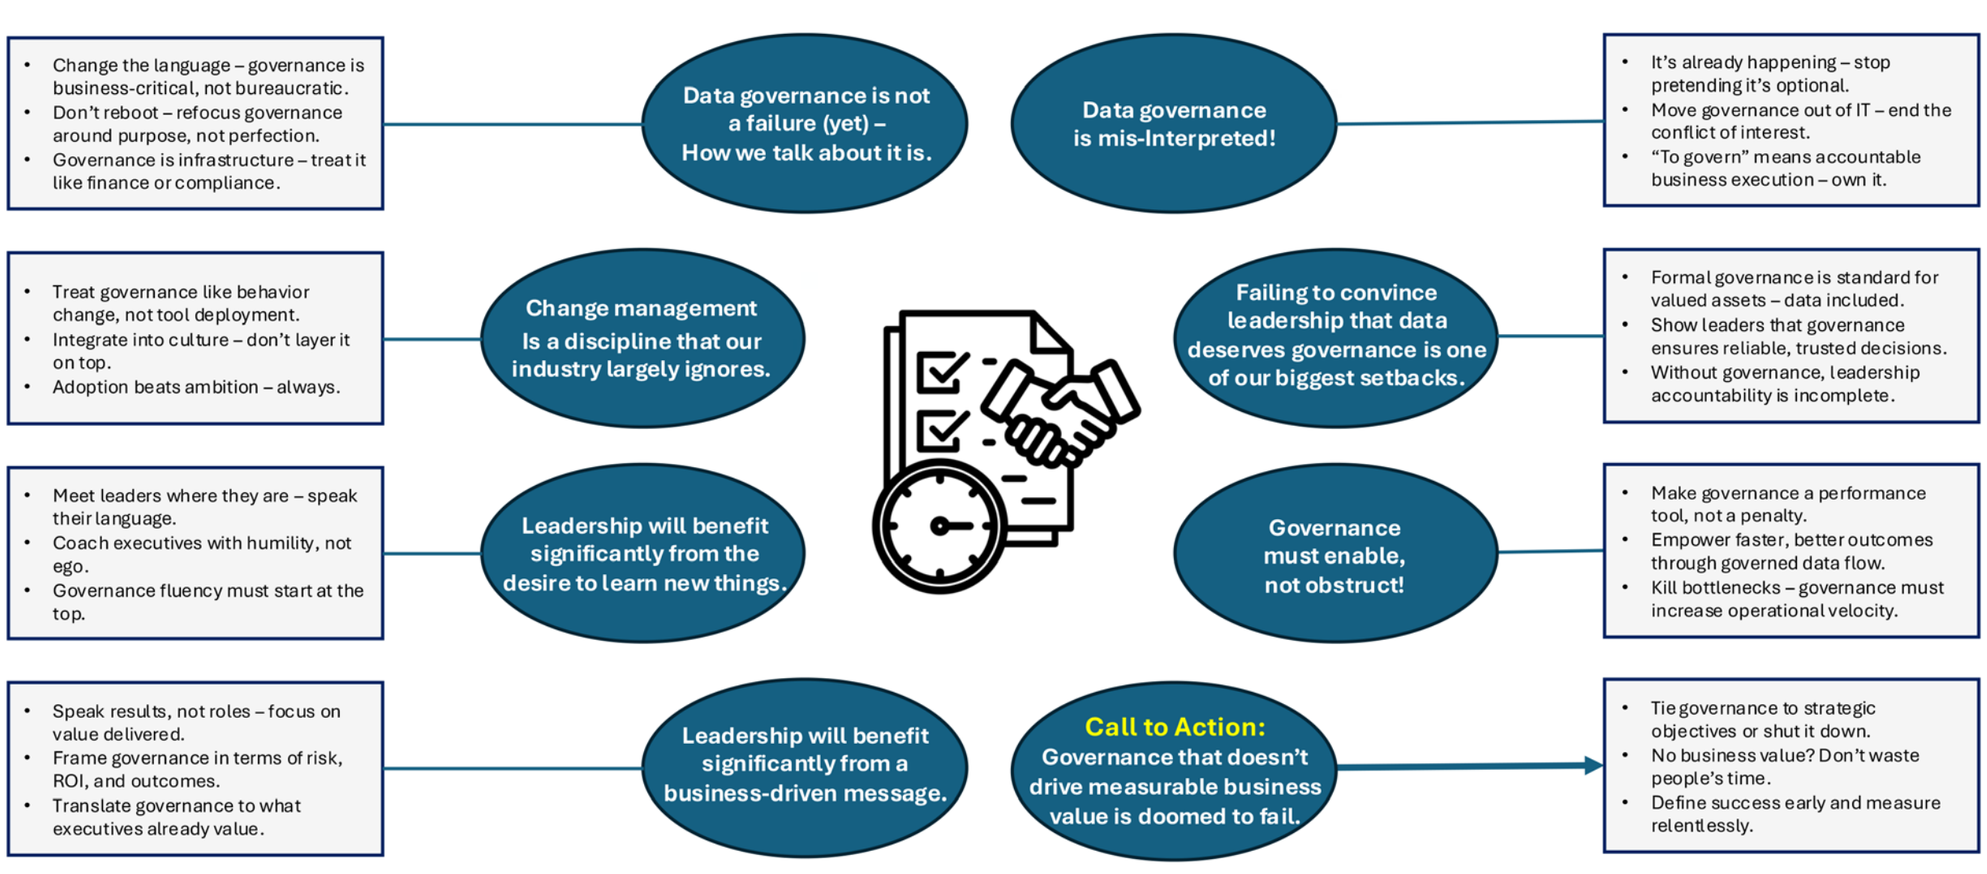

Guide to PDF ingestion for AI

A detailed guide clarifying how AI systems should process PDF documents has been published by The PDF Association. The FAQ warns that common ingestion practices risk information loss and hallucinations.

It offers practical guidance for information managers preparing document collections for AI ingestion.

The resource argues PDF content is highly valued by AI systems because PDFs function as the established "document of record" in human communication. This contrasts with HTML web pages, which the FAQ describes as transactional, short, and subject to change.

Citing a recent HuggingFace blog post, the Association claims PDF content offers higher information density and is inherently long-context.

A central warning concerns converting PDFs to plain text or Markdown before ingestion. The FAQ calls this approach "inevitably lossy" and an unnecessary "dumbing down" process that risks increasing AI hallucinations.

Stripping semantic information such as strikethrough, superscript, or table structure removes intended meaning. The FAQ uses the example of plain "22" versus a superscripted "2²" to illustrate how conversion loses mathematical, chemical, or footnote context.

The FAQ argues AI systems should also avoid processing PDFs page by page, because content routinely breaks across page boundaries. Isolating individual pages reduces understanding and raises hallucination risk.

Tagged PDF as key enabler

For records managers and digital transformation teams, the FAQ places particular emphasis on Tagged PDF. Tagged documents provide logical reading order, natural language indicators, table structure, and alt-text for images.

These tags give AI systems an unpaginated logical structure, avoiding the need to interpret pagination artefacts. The Association recommends born-digital documents conforming to WTPDF, PDF/UA-1, or PDF/UA-2.

Tables receive specific attention. Because PDF does not natively define tables in its graphical model, AI systems processing untagged PDFs must rely on error-prone visual analysis. Tagged PDF tables, by contrast, express semantics that align closely with HTML.

AI ingestion systems should capture all PDF components, including annotations, embedded files, bookmarks, and XMP metadata. Ignoring this information, it argues, reduces understanding and increases hallucination risk.

The FAQ cautions that PDFs containing unresolved redaction annotations may expose personally identifiable information during AI ingestion. Such annotations flag content for removal but do not purge it.

The guide also addresses security permissions. Some AI systems ignore author-set restrictions indicating that text and graphics should not be reused.

The PDF Association is the industry body representing PDF technology vendors and promotes ISO-standardised PDF features.

Read the full FAQ at https://pdfa.org/faq-ai-and-pdf/

OpenAI releases workspace agents

OpenAI has introduced workspace agents in ChatGPT, enabling business teams to build shared AI agents that automate complex, long-running workflows across an organisation.

The feature is powered by OpenAI's Codex model and represents a significant evolution of the company's earlier custom GPT functionality. Unlike custom GPTs, which responded to individual conversational prompts, workspace agents run continuously in the cloud and can execute multi-step tasks without repeated human input.

Teams can use workspace agents to prepare reports, write code, respond to messages, and route information across business tools including Slack. Agents can be built once and shared across an organisation, with the ability to retain context between sessions and connect to external applications.

"Workspace agents are designed for that kind of work," OpenAI said in its announcement. "They can gather context from the right systems, follow team processes, ask for approval when needed, and keep work moving across tools."

The capability is available as a research preview for ChatGPT Business, Enterprise, Edu, and Teachers subscribers. Enterprise and Edu administrators can activate the feature through role-based access controls that govern which users can create agents and which tools they are permitted to use.

Workspace agents are free to use until 6 May 2026, after which OpenAI will apply a credit-based pricing model. The company did not disclose specific pricing details at launch.

OpenAI said it is taking steps to protect workspace agents from prompt injection attacks and other security threats. Organisations can restrict the data and tools agents are permitted to access, and can require human approval before an agent executes sensitive actions.

Example use cases outlined by OpenAI include a software review agent that triages IT requests and enforces procurement policy, a product feedback routing agent that draws from Slack and support channels, and a weekly metrics reporting agent that automatically pulls data, generates charts and delivers business reports.

Existing custom GPTs will remain available while organisations evaluate workspace agents. OpenAI said it plans to release a migration tool that will allow users to convert their custom GPTs into workspace agents at a later date.

The launch intensifies competition in the enterprise agentic AI market, where Google, Microsoft, and Amazon are each investing heavily in autonomous workflow tools. OpenAI framed workspace agents as an early step toward AI systems that can collaborate alongside human workers within the tools and processes where work already occurs.

Workspace agents can be created by clicking the Agents option in the ChatGPT sidebar and describing a workflow in natural language. ChatGPT then guides the user through configuring steps, connecting tools, and testing the agent before deployment.

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

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

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

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

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

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

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

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

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

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

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| 1300 474 288

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

Informotion acquires InTouch Suite

Informotion has acquired the InTouch product suite from InReach Software, taking ownership of a long-running web and mobile interface designed to extend OpenText Content Manager by enhancing the user experience and increasing productivity gains.

The acquisition consolidates an existing arrangement under which Informotion already resold and implemented InReach Software solutions for Content Manager customers. The deal formalises that relationship by bringing the product into Informotion's own portfolio.

The full InTouch product suite includes the InPort integration engine and Export manager. InTouch helps automate existing processes and workflows while offering a way to electronically approve and endorse items. Export Manager provides custom report metrics out of both Content Manager, and inTouch, for use with an organisation's PowerBI team. InPort helps onboard staff, create and manage locations, and automate folder creation including Personnel File folder structures.

InTouch sits on top of an existing Content Manager environment, either on-premise or Content Manager Cloud, without requiring new infrastructure, desktop rollout or changes to system configuration. The change is at the user experience layer: a browser-based interface intended to make everyday tasks such as document creation, approvals, correspondence and records capture easier on any device.

Informotion is positioning the acquisition as a response to a well-known issue in the OpenText Content Manager community: powerful records management capability that does not always reach end users, leading to inconsistent record capture, workflow non-completion and quiet compliance drift.

"Most Content Manager customers tell us that while the platform is incredibly powerful, much of that capability never reaches the hands of their users," said Shane Parsons, CEO of Informotion.

"The result is friction, especially for occasional users or staff working across devices. Records don't get captured. Workflows don't get followed. Compliance quietly suffers in the background."

Through a clean, modern, browser-based interface, users can access Content Manager from any device, making everyday tasks such as document creation, approvals, correspondence, and records capture simple, fast, and intuitive. This is particularly valuable for government document generation and approval workflows, where ease of use, process adherence, and auditability are critical.

Beyond the interface, InTouch introduces automated workflow and process management, enabling organisations to digitise and streamline approval and correspondence processes with end-to-end audit trails. A built-in Management Dashboard provides leaders with real-time visibility across task status, workload distribution, and bottlenecks supporting better decision-making and governance.

"With this acquisition, we see a real opportunity to help Content Manager customers modernise the user experience, drive adoption, and strengthen compliance outcomes - all while maximising their existing investment," Parsons said.

"The IP for the InTouch Product suite became available to purchase. With our deep product expertise, along with our many existing customers with the knowledge of the benefits the product brings to our clients, including increasing their use, efficiency and compliance of Content Manager, it made perfect sense for us to take this step and continue the journey."

The deal is significant for the Australian and New Zealand records management community given the position of OpenText Content Manager, formerly TRIM and HP Records Manager, as the dominant electronic document and records management system in Australian and New Zealand government since the late 1980s. The platform remains in heavy use across federal, state and local government, regulated industry and the corporate sector.

InTouch not only provides a simple and clean interface for accessing data within Content Manager, it also helps automate and ease common pain-points such as automated titling, automated templates, automated folder creation, and structured workflow with built in document previewer and secure editing within Microsoft 365.

InTouch workflows plug the gaps in current Content Manager Action Tracking and Workflow while keeping things simple for users and ensuring sensitive information is kept secure. For Ministerial briefings where the process might require portions of a document to be allocated to two teams at the same time for a consolidated response, InTouch provides this capability

Informotion supports Content Manager customers across Australia, New Zealand, the Europe Middle East and Africa region and the United Kingdom, and is an established OpenText partner with delivery teams in Sydney, Canberra, Melbourne, Adelaide, Perth, Brisbane and the UK.

https://www.informotion.com.au

Hyland expands AI Content Services

Hyland has released a set of AI-driven product updates targeting accounts payable, SAP S/4HANA, and Workday Student integrations, alongside platform-wide enhancements to its Content Innovation Cloud.

The announcements centre on embedding AI agents directly into existing business systems rather than deploying standalone tools, a shift the company frames as moving from AI that suggests to AI that can act within governed processes.

Hyland Agentic General Ledger Coding replaces manual GL coding in accounts payable workflows

with AI-generated recommendations. The system provides confidence-scored predictions and retains human-in-the-loop oversight.

The company says it supports compatibility with a range of enterprise resource planning (ERP) systems.

Hyland for SAP S/4HANA is positioned as a clean-core integration, embedding content management and process automation directly into SAP's Fiori interface without custom code. The company says this supports auditability, compliance, and upgradability in SAP environments.

Hyland for Workday Student extends document management into Workday Student workflows, covering upload, search, and management of student records. The integration is described as cloud-native and designed to prepare student content for use with AI applications.

Platform-wide updates include enhancements to the Hyland Intelligent Document Processing (IDP) tool, the Knowledge Discovery search service, and the Knowledge Enrichment capability for converting unstructured content into structured, AI-ready data.

The Content Federation Service (CFS) is now offered within the Content Innovation Cloud, connecting OnBase, Alfresco, Nuxeo, Perceptive Content, and SharePoint 365 to cloud-native AI capabilities without content migration.

"Organizations don't need more disconnected AI tools; they need intelligence embedded where work actually happens," said Michael Campbell, chief product officer at Hyland.

"Our latest product innovations extend content management and AI directly into the systems our customers rely on every day, turning insight into action while preserving the governance, trust, and control required to operate at enterprise scale."

https://www.hyland.com

Alation Launches AI Governance Platform

Enterprises rolling out artificial intelligence faster than they can govern it now have a new option from data catalogue vendor Alation, which has introduced Alation AI Governance, a system of record for AI models, agents and tools. The offering registers every AI asset in a single inventory, maps each to applicable regulations, generates evidence-backed model cards, routes approvals through regulation-aware workflows and produces a live compliance posture for executives.

Alation says the product addresses a gap created when organisations deploy AI faster than governance frameworks can be applied. The vendor describes a familiar scenario in which model documentation lives in SharePoint pages, approval threads sit in email and stale records cannot satisfy regulators or board questions. The product registry covers built-in support for the EU AI Act, AI-relevant subsets of the General Data Protection Regulation, the United States National Institute of Standards and Technology AI Risk Management Framework, and ISO 42001. Customers can add additional regulations, with AI-assisted suggestions to accelerate requirement mapping.

Five capabilities sit inside the product. An AI Asset Registry ingests assets from connected platforms or via software development kit submission. AI-Native Model Cards are generated from metadata, data dependencies and regulatory requirements.

The Executive Dashboard presents an overall compliance score, per-regulation breakdown with trend lines and the top open risk items blocking compliance, with drill-down to underlying assets.

https://www.alation.com

Agent Activity Gains New Visibility

Engineering teams running AI agents in production have a new way to see what those agents are actually doing, after observability vendor Honeycomb introduced a set of agent-focused features built on OpenTelemetry standards rather than proprietary SDKs.

The release adds three capabilities to the Honeycomb platform: Agent Timeline, Canvas Agent and Canvas Skills. The vendor frames the launch as a response to the non-deterministic, multi-hop nature of agent workflows, which it argues older observability tools were not designed to handle.

Agent Timeline renders multi-agent, multi-trace workflows as a single view, connecting language model calls, tool invocations, agent handoffs and downstream system impacts. The intent is to let engineers reconstruct a decision path when an agent contributes to an incident, rather than piecing together logs by hand.

Canvas has been rebuilt as a collaborative workspace, chat interface and autonomous agent in one. Auto-investigations let the Canvas agent gather data, form and test hypotheses and propose remediation when an alert fires or a service level objective is at risk. Canvas Skills encode debugging knowledge and framework-specific practices, such as Kafka troubleshooting, into reusable playbooks.

The product approach leans on OpenTelemetry rather than custom instrumentation. Honeycomb has integrated the OpenTelemetry GenAI semantic conventions at version 1.40.0, treating gen_ai.* attributes as first-class data so model evaluations, tool executions, Model Context Protocol calls and language model outputs surface without re-instrumentation as the specification evolves.

https://www.honeycomb.io

SAP Bets ERP Future On Agentic AI

SAP has used its Sapphire conference in Orlando to position itself as a business AI company rather than an enterprise software vendor, unveiling what it calls the Autonomous Enterprise alongside a unified AI platform, a portfolio of more than 200 agents and a €100 million fund for partner-led deployments.

The launch combines three elements. SAP Business AI Platform brings together SAP Business Technology Platform, SAP Business Data Cloud and SAP Business AI into a single governed environment. SAP Autonomous Suite layers AI agents over existing applications to run processes end-to-end. Joule Work becomes the front-door user experience across desktop, mobile and voice.

More than 50 domain-specific Joule Assistants will sit across finance, supply chain, procurement, human capital management and customer experience.

SAP says these will orchestrate a subset of over 200 specialised agents. As one illustration, an Autonomous Close Assistant is intended to compress the financial close cycle from weeks to days through automated journal entries, reconciliation and error resolution.

An Industry AI portfolio adds seven autonomous solutions that embed sector-specific process logic, data models and regulatory requirements. SAP showcased work with German utility RWE under an Autonomous Asset Management scenario, in which agents analyse incident data to identify likely root causes and generate pre-filled work orders for offshore wind turbines.

At the centre is the SAP Knowledge Graph, which provides AI agents with a structured map of business entities, processes and relationships across a customer's SAP landscape. Joule Studio is positioned as SAP's AI-first build environment for enterprise agents, applications and workflows, with no-code, pro-code and AI framework options.

"For the mission-critical processes of our customers, 'almost right' just isn't good enough," said Christian Klein, CEO of SAP SE. "By uniting SAP Business AI Platform with SAP Autonomous Suite, we anchor AI agents in the business processes, data and governance so they can deliver accurate, compliant and secure outcomes, unlocking new sources of revenue and meaningful cost savings."

SAP also announced a wide partner slate. Anthropic's Claude is among the foundation models powering Joule agents across human resources, procurement and supply chain. Amazon Web Services will provide zero-copy data integration between SAP Business Data Cloud and Amazon Athena. Google Cloud and Microsoft enable bidirectional agent-to-agent interoperability between Joule and external agent frameworks.

Mistral AI and Cohere supply sovereign model options on SAP's cloud infrastructure. NVIDIA's OpenShell underpins the Joule Studio runtime, isolating each agent in a sandboxed environment with configurable policies. Palantir and Accenture are positioned as implementation partners for complex data migration scenarios.

Customer transition pathways have been reset. RISE with SAP customers receive a contractual commitment to activate three Joule Assistants within their first year, while SAP GROW customers get full portfolio access at onboarding. SAP S/4HANA on-premises and ERP Central Component customers that commit to transitioning the majority of their landscape to SAP Cloud ERP gain access to selected AI scenarios.

SAP also said new agent-led transformation tooling could reduce ERP migration effort by more than 35 percent through automated system analysis, code remediation, configuration and testing.

Oracle, Microsoft, Workday and Salesforce have all announced overlapping agentic capabilities for finance, supply chain and human capital management.

https://www.sap.com

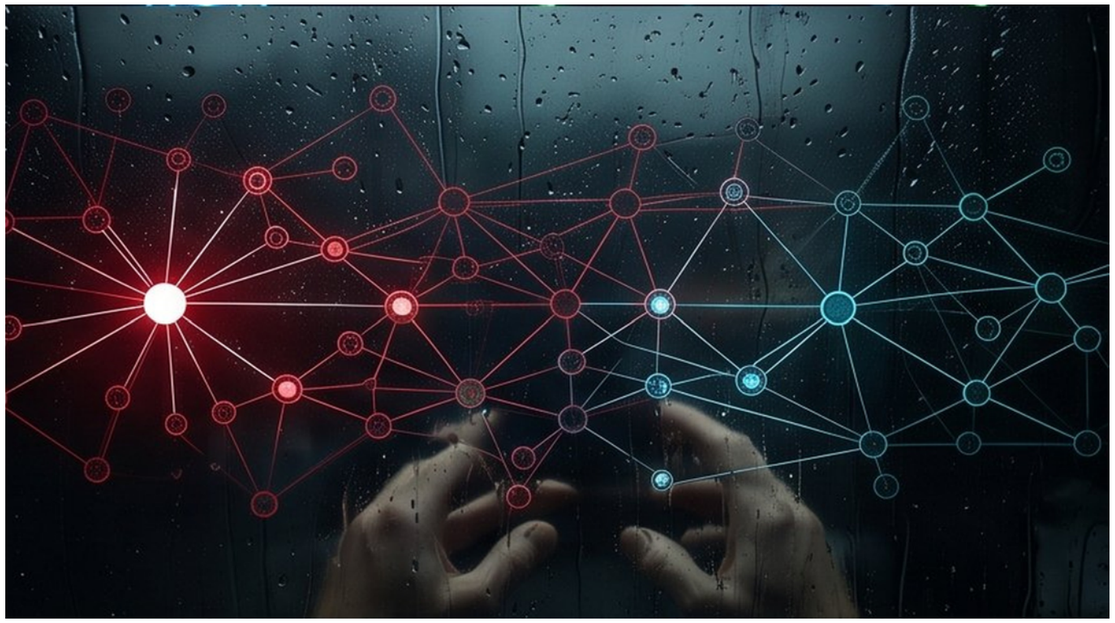

Rossum IDP Joins Coupa Stack

Spend management vendor Coupa has acquired intelligent document processing specialist Rossum, moving a two-year technology partnership inside its own walls and extending the capability across its full source-to-pay platform.

The two companies have worked together since 2024 on automating complex invoicing for accounts payable teams, with Rossum's document processing already embedded in the Coupa platform for that workflow.

Coupa says the next step is to take Rossum's capability beyond accounts payable and into the broader Coupa portfolio, covering both direct and indirect spend. The intent is to combine Rossum's transactional intelligence with Coupa's Navi agentic AI fleet across source-to-pay.

Rossum positions its architecture as moving beyond legacy optical character recognition. Its transactional intelligence is powered by what the company calls a transactional large language model, or T-LLM, trained on tens of millions of documents and described as continuously learning from each customer's document set.

Tomáš Gogár, Rossum CEO and Co-Founder, framed the deal as the natural evolution of a partnership built on a shared AI-first culture.

"By combining our proprietary T-LLM transactional intelligence with Coupa's massive $10T data set, we are well positioned to create immediate customer value and fundamentally change how the world buys and sells."

https://www.coupa.com

Affinda agent handles IDP configuration

Document automation vendor Affinda has launched an AI agent that builds end-to-end document workflows through a single natural language conversation. The Melbourne-headquartered firm claims it can compress configuration cycles that previously took six to nine months.

The Affinda Agent configures ingestion, splitting, classification, extraction, validation and integration steps from a plain-language description of the user's documents, business logic and target systems.

Users can inspect, adjust or reject each proposed configuration before it takes effect. The agent runs inside the existing Affinda Platform and connects to more than 2,800 business systems.

Affinda claims its platform is now the only intelligent document processing (IDP) solution that configures an entire document workflow through a single natural language conversation.

Andrew Bird, Head of AI at Affinda said: "For too long, the barrier to document automation hasn't been the AI - it's been the setup. Every other approach still requires technical teams to stitch together the workflow: ingestion, splitting, classification, extraction, validation, integration - each piece configured separately. We've collapsed that entire process into one conversation. That's not an incremental improvement; it's a different category of product."

Last year Affinda launched a new platform deploying AI agents with persistent Model Memory, enabling custom workflows without lengthy training cycles.

https://www.affinda.com

Semantic search expands Appian AI

Two new AI capabilities - Agent Studio and AI Document Centre - have reached significant milestones on the Appian Platform, with implications for organisations managing document-intensive compliance and operational workflows.

Agent Studio, now in beta, enables the design and deployment of AI agents capable of executing multi-step tasks, interacting with multiple systems, and making decisions based on real-time data and business logic.

The platform allows agents to update records, send emails, and respond dynamically to changing inputs - functionality relevant to records managers and GRC teams handling complex approval chains or regulatory submissions.

AI Document Centre has reached general availability. The tool is designed for enterprise-scale intelligent document processing, handling complex formats with what Appian describes as high extraction accuracy. Appian customer Century Fire Protection reported a 36% reduction in invoice operating time after using the tool to automate accounts payable - one of the few concrete performance figures in the announcement.

Additional platform enhancements include semantic search across the Appian data fabric, including text fields and attached documents. Smart search interprets intent and surfaces related records beyond keyword matching - a capability with practical applications for information retrieval in legal, compliance, and case management contexts.

The release also introduces autoscale support for generative AI agents in high-volume process environments. Appian describes the feature as delivering 10 to 100 times the throughput without over-provisioning.

Data fabric enhancements include native document management within record types, incremental external data sync as frequently as every 15 minutes, and support for up to 20 million rows per record. Form design improvements include automated header and wizard functionality.

https://www.appian.com

Security Firm Targets Credential Theft

Detecting credential theft as it happens is the aim of Decipio, a new community security tool from Arctic Wolf.

The tool targets a well-documented Windows attack vector, aiming to alert defenders before stolen credentials are weaponised inside a network.

Decipio operates as a passive network tripwire. It monitors for exploitation of LLMNR (Link-Local Multicast Name Resolution) and NBT-NS (NetBIOS Name Service) - two legacy Windows name-resolution protocols still enabled by default in all current Windows versions.

When an attacker on the same network segment poisons these protocols, victim machines send authentication credentials directly to the attacker. Arctic Wolf says Decipio is designed to generate a binary alert at this moment, requiring minimal tuning. These performance claims could not be independently confirmed.

The technique - catalogued by MITRE ATT&CK as sub-technique T1557.001 - has been a standard tool in attackers' arsenals since at least 2012. Both LLMNR and NBT-NS remain enabled by default in Windows 11 and Windows Server 2025.

The US Cybersecurity and Infrastructure Security Agency (CISA) published a formal countermeasure in March 2025 recommending LLMNR be disabled entirely. Microsoft is separately pursuing a phased roadmap to deprecate NTLM - the authentication protocol underpinning these attacks - by default in future Windows releases.

https://arcticwolf.com/decipio/

Data Quality Scores Go Live in ServiceNow

Poor data quality is a leading reason AI initiatives stall in production. Enterprise teams using ServiceNow will now see data quality scores directly on assets in its Data Catalogue, before those assets power AI models or business decisions.

Data management software provider Ataccama has announced an integration with ServiceNow that surfaces quality assessments alongside data assets in the ServiceNow Data Catalogue.

These include colour-coded quality-tier badges, specific failure findings, and observability signals such as schema changes and volume anomalies.

When a ServiceNow user or AI agent selects a data asset, it can immediately view Ataccama's quality assessment. Data lineage, downstream impact and business criticality information are also visible within the same interface. Teams can trace root causes and prioritise fixes without switching tools.

The two companies say the integration addresses a core challenge in enterprise AI deployment. Without quality signals embedded at the point of data discovery, AI workflows and automated decisions may proceed on data that has not been validated for fitness.

Quality scores are based on Ataccama's Data Trust Index. Data stewards and business users in ServiceNow work from the same quality signals monitored by data engineers in Ataccama, according to the company.

A shared quality view aims to reduce reconciliation across tools and surface issues before they affect downstream outcomes.

https://www.ataccama.com

Agentic AI Bundles for the CFO and CIO

Automation Anywhere has launched two pre-built agentic AI suites targeting whole IT and finance functions. The company pitches them as a route past the stalled pilot phase plaguing much enterprise AI investment. Autonomous IT and Autonomous Finance combine AI agents, process intelligence, key performance indicators, governance controls and pre-built connectors for major enterprise platforms.

Autonomous IT is aimed at the office of the CIO. It covers IT service management, cloud operations, security operations, compliance, identity and access management, financial operations and service continuity. The package ships with more than 45 agents handling over 70 tasks.

Autonomous Finance covers quote-to-cash, procure-to-pay, record-to-report, financial planning and analysis, treasury, payroll, tax and controls. It includes more than 55 agents covering 100-plus tasks, with built-in enterprise resource planning integrations. Automation Anywhere claims customers can achieve more than 90 per cent straight-through processing for accounts payable using the finance package.

"Enterprises do not need more isolated AI experiments. They need a practical path to run core functions with greater autonomy, control, and measurable value," said Mihir Shukla, CEO and Board Chairman of Automation Anywhere.

"Autonomous IT and Autonomous Finance give leaders a faster starting point with pre-built agents, process intelligence, controls, and roadmaps designed for how their functions actually run."

https://www.automationanywhere.com/

Enterprise Vault Merges Tools

Enterprise Vault has consolidated its archiving, eDiscovery, surveillance and data insight tools into a single bundled offering.

The package targets organisations that need to keep regulated data inside their own environments.

Enterprise Vault Complete pairs the Merge1 communications capture connectors with the vendor's existing archiving, search, supervision and content classification modules under one licence.

The bundle targets compliance teams managing fragmented governance tooling.

Core features include:

- Data archiving, capture and classification capabilities that support retention, preservation and policy-based governance
- eDiscovery capabilities that help legal and compliance teams search, review and investigate with greater efficiency
- Surveillance and data insight capabilities that strengthen oversight across communications and data activity
- Merge1 connectors that expand capture across modern communications sources

Enterprise Vault now operates as a business unit of Cloud Software Group.

The product line itself has a long lineage running through Veritas, Symantec and most recently Arctera.

https://www.enterprisevault.com/

Sensitive Data Guardrails in Purview

Microsoft is expanding its Purview Data Loss Prevention platform to intercept Microsoft 365 Copilot prompts that contain sensitive data, blocking AI-generated responses before they are returned to users.

The capability, currently rolling out as part of Microsoft's 2026 product roadmap, prevents Copilot - including pre-built agents in Microsoft 365 Copilot - from processing or grounding responses on content that matches DLP policy conditions.

According to Microsoft's product documentation, the control is designed to mitigate data leakage and oversharing risks that arise when employees interact with AI assistants using company data.

When a prompt containing sensitive information is detected, Copilot is prevented from returning a response.

The update is one of several Purview enhancements rolling out in 2026. A companion feature, Insider Risk Management for risky AI usage, adds detection of intentional and unintentional insider risk activity across generative AI applications - including third-party platforms beyond Microsoft's own products.

Microsoft has also extended Purview data classification across SharePoint, Teams and OneDrive, and introduced improved compliance reporting for AI-generated content.

For organisations that have deployed or are evaluating Microsoft 365 Copilot, the Purview controls address a governance gap that has concerned compliance and records management teams.

Copilot's ability to surface content from across a Microsoft 365 environment has raised concerns about whether sensitive or classified content could be inadvertently exposed through AI-generated responses.

A Gartner survey of IT leaders found data oversharing caused 40 per cent of organisations to delay Copilot deployments by three months or more.

Docusign AI Assistant for Contract Reviews

Organisations using Docusign's Intelligent Agreement Management platform can now deploy an AI-powered contract review assistant that analyses agreements, flags risky terms, suggests edits, and generates legal playbooks.

The tool - called AI-Assisted Review - is powered by Iris, Docusign's proprietary AI engine. It lets users interrogate contract terms in plain language, asking questions such as whether a contract auto-renews, and receive answers linked to specific clauses.

The assistant also supports playbook management. Legal teams can compare contracts against existing playbooks to identify deviations, or generate new playbooks from standard templates.

Docusign says this standardises review across teams, enabling staff outside the legal department to handle routine agreements with less central oversight.

The feature is available to Docusign CLM and Agreement Desk customers on select plans, globally.

"Contract review has traditionally been one of the biggest bottlenecks in how agreements move through a business," said Dmitri Krakovsky, Chief Product Officer at Docusign.

"With the new contract review assistant, legal teams can quickly understand contracts, identify risks, and generate edits in seconds. Because it's built on the Docusign IAM platform, those actions flow directly into the broader agreement workflow, helping legal stay connected with sales, procurement, HR, and other teams across the company."

Docusign cited internal performance data claiming its own legal team saved up to 15 minutes per NDA review and reduced MSA negotiation time by 30 to 60 minutes.

The company also cited a Deloitte 2025 report, in which over 70 per cent of legal leaders said agreement management tools improved caseloads and legal outcomes.

The launch comes as AI adoption for contract review is accelerating rapidly across in-house legal teams.

A LegalOn Technologies survey of 452 in-house legal professionals, published in January 2026, found that active AI use in contract review had nearly quadrupled since 2024, with more than half of teams using or evaluating the technology.

AI-Assisted Review is integrated with the broader Docusign IAM platform, allowing reviewers to move between IAM and Microsoft Word within agreement creation workflows. Docusign has not detailed support for other document environments.

https://www.docusign.com

Centralising AI Agent Orchestration

Automation Anywhere has unveiled new platform capabilities for orchestrating AI agents, automations and human workers within a single governed enterprise process.

The release targets organisations moving AI from pilots into production operations. CEO Mihir Shukla said the autonomous enterprise requires a system that coordinates how work runs within departments and across the organisation.

New universal orchestration capabilities sequence tasks across departments and platforms including Salesforce, ServiceNow and SAP. The vendor said the orchestration layer handles approvals, data sources, APIs, automations, AI agents and human decisions in one process.

Automation Anywhere also introduced Automation

Anywhere Code (AAI Code), a low-code tool for building enterprise applications. The vendor said AAI Code can produce apps complete with UI, processes, agents, context and security in as little as one week.

The company described AAI Code as distinct from vibe coding tools because it plans before building and includes governance controls before deployment.

A new Context Intelligence Graph sits within the Process Reasoning Engine. It is designed to surface relevant context for each task rather than exposing all enterprise data to every step.

Automation Anywhere said agents using the engine demonstrated more than 30 per cent higher accuracy in internal evaluations across multiple business domains.

For governance and risk teams, the additions include AI Evaluations to assess agent performance at design and runtime. A new Process Simulation, Optimization and Testing environment stress-tests entire processes including exceptions and edge cases before deployment.

The vendor cited University Hospitals of Leicester NHS Trust as targeting 50 to 70 per cent autonomy for administrative work. Anticipated benefits include a 22-day reduction in recruitment time and a one million pound annual cut in temporary staffing costs.

https://www.automationanywhere.com

Sovereign AI needs workflow control

As enterprises accelerate deployment of agentic AI systems, Automation Anywhere has published a framework arguing that conventional sovereign AI approaches - focused primarily on data storage location - are insufficient for workflows where AI agents actively execute tasks across multiple systems and jurisdictions.

The company's "spectrum of control" model defines sovereignty across five dimensions: where data and metadata reside; how data is processed and whether it is copied or moved; who controls access and encryption keys; where actions are performed; and which legal jurisdictions govern data access.

The model is positioned as a response to enterprise environments where workflows routinely span multiple clouds, on-premises infrastructure, and regulatory regimes.

According to McKinsey research cited by Automation Anywhere, three-quarters of countries have implemented data localisation rules, creating compliance complexity for global enterprises attempting to standardise AI operations.

For organisations in sectors such as financial services, healthcare, or government, agentic AI that moves data across systems to complete tasks introduces new audit and liability exposure that data residency controls alone do not address. (

Automation Anywhere's Agentic Process Automation (APA) platform is cited as supporting flexible deployment across cloud, multi-cloud, and on-premises environments without requiring data centralisation.

Key governance capabilities include composable architecture avoiding vendor lock-in, action and workflow controls defining how AI agents operate across environments, and audit trail requirements for compliance and responsible AI obligations.

https://www.automationanywhere.com

NetDocuments Bridges the Document Gap

Legal teams managing complex litigation portfolios have long operated across disconnected systems - switching between case management platforms and document repositories to track obligations, deadlines and files. A new integration between cloud-based case management platform Clarra and document management system (DMS) provider NetDocuments aims to eliminate that switching.

The partnership allows legal teams to access and manage NetDocuments workspaces and files directly from within the Clarra platform. Users can create a matter in Clarra and automatically generate a corresponding NetDocuments folder structure or link existing NetDocuments repositories to active matters.

Document repositories connect to Clarra matters, keeping case data, deadlines and obligations aligned in a single interface.

The integration preserves the NetDocuments interface for users who are already familiar with it, which the companies say reduces change management and training requirements.

Live document repositories tied to matters are intended to reduce duplicate files and version control errors.

Clarra's integration model reflects a broader strategy of supporting legal technology stacks assembled from specialist best-of-breed tools, rather than requiring adoption of all-in-one platforms. The company says it integrates with more than 6,000 applications.

https://www.clarra.com

Cloudera targets Data modernisation

Cloudera has released a set of updates to its hybrid data and AI platform that extend vendor support through 2032, introduce on-demand cloud bursting, and add automated query optimisation for Apache Iceberg table formats.

The updates are designed to help enterprises modernise data infrastructure without disruptive migration projects, while also reducing infrastructure costs and improving performance across cloud and on-premises environments.

A key feature is the Cloudera Lakehouse Optimizer, which automates the maintenance of Apache Iceberg tables. Cloudera claims the tool accelerates query performance by 38% and reduces storage overhead by up to 36% with minimal manual effort.

Cloudera Cloud Bursting allows organisations to dynamically extend on-premises workloads into public cloud environments on demand, without requiring data duplication or application rewrites. The company says this enables elastic scale while protecting existing data centre investments.

Expanded data sharing capabilities allow live Apache Iceberg tables to be accessed across external platforms without copying data. Cloudera says this reduces data silos while maintaining governance and data integrity.

Extended platform support through 2032 is positioned as a mechanism for reducing upgrade risk and aligning infrastructure investment with longer planning cycles, particularly relevant to large enterprises in regulated sectors.

Apache Iceberg has emerged as the dominant open table format for large-scale analytics workloads.

https://www.cloudera.com

AI platform for underwriting and claims

Duck Creek Technologies has launched an insurance-specific agentic AI platform designed to deploy, orchestrate, and govern AI agents across property and casualty workflows - targeting what Boston Consulting Group estimates could represent up to $80 billion in annual industry impact in the United States alone.

The Duck Creek Agentic AI Platform combines generative AI, machine learning, and neuro-symbolic reasoning - a hybrid approach pairing probabilistic AI with deterministic rules-based logic. The company argues this combination is essential for insurance workflows where decisions must satisfy regulatory requirements and be fully explainable.

The platform includes a proprietary insurance-domain Model Context Repository (MCR), an orchestration layer for designing and deploying agents, and a built-in governance module providing decision traceability, auditability, and compliance controls.

Two initial agentic applications accompany the platform launch. The Agentic Underwriting Workbench automates submission intake, triage, and enrichment, prioritising high-value opportunities and preparing decision-ready submissions to reduce quote turnaround times.

The Agentic First Notice of Loss (FNOL) application handles claims intake across digital, voice, and mobile channels, applying AI to policy verification, data validation, and early fraud detection. The FNOL application was developed in collaboration with Google Cloud using Gemini models.

The platform is structured in five layers: agentic intelligence (the insurance domain model); orchestration (agent design and deployment); AI Assurance (governance and auditability); AI Gateway (ecosystem integration including MCP and A2A protocols); and native integration with Duck Creek's core policy, billing, and claims systems.

For insurers already running Duck Creek core systems, the platform leverages existing data, configurations, and APIs rather than requiring data migration.

https://www.duckcreek.com

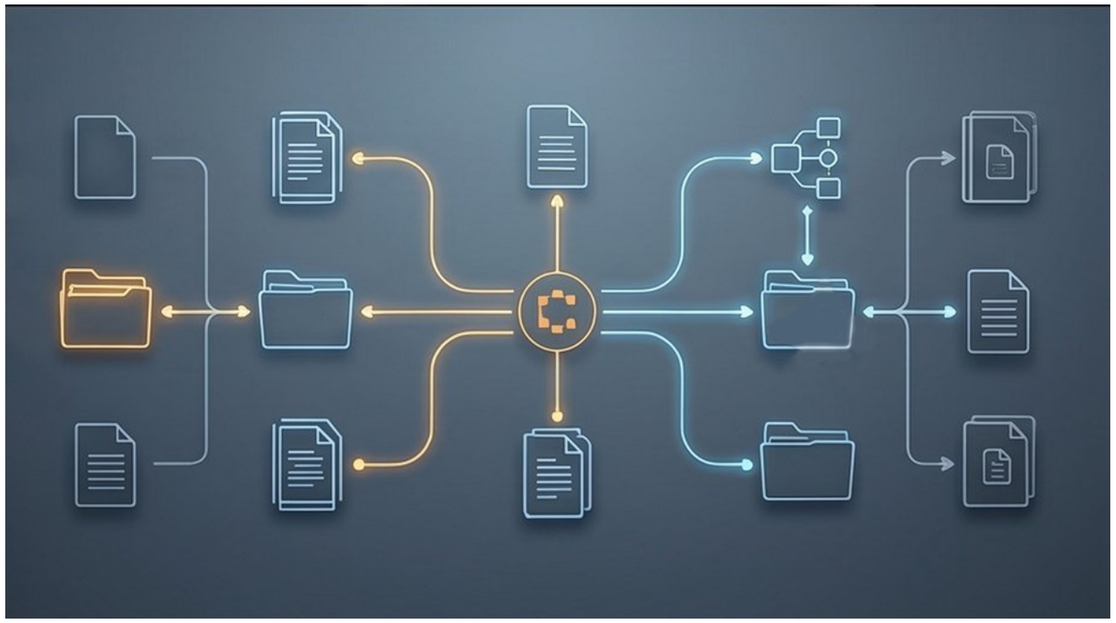

Pure Storage rebrands as Everpure

Pure Storage has rebranded as Everpure and completed its acquisition of data intelligence vendor 1touch.

The company describes the moves as a pivot from storage hardware specialist to broader data management platform.

1touch provides data discovery, classification and contextualisation across SaaS, on-premise and edge environments.

Everpure says the technology will fold into its Enterprise Data Cloud architecture. The goal is to make data discoverable and labelled at the storage layer. Embedding classification into the storage platform could shift workloads currently handled by separate data catalogs, sensitive data tools or content services platforms.

"With 1touch, we are taking the next step in helping organisations not only gain control of their most valuable asset – data - but also understand, enhance, and contextualise that data for actionable intelligence," said Everpure CEO Charles Giancarlo.

He said data trapped in silos and inflexible architectures cannot support the scale and speed of enterprise AI deployments.

The company is pitching Enterprise Data Cloud as a single virtualised layer over distributed storage, governed by an intelligent control plane.

"Data is the lifeblood of the AI era, but without the proper controls and semantic context, it remains an untapped resource," said Ashish Gupta, CEO and president, 1touch."

https://www.everpuredata.com/

Flexor AI Preps Unstructured Data

A persistent barrier to production-grade enterprise AI - that agents lack sufficient business context to deliver consistent results - is the target of a new integration between Flexor's AI Context Engine (ACE) and the Snowflake AI Data Cloud.

The integration processes unstructured enterprise content - emails, call transcripts, PDFs, support tickets, and notes - and transforms it into structured, AI-ready context that operates within Snowflake without moving data to external systems.

The core problem ACE addresses is well-documented across enterprise AI deployments: large language models and AI agents produce unreliable outputs when they lack organisational knowledge - company-specific terminology, naming conventions, relationships between documents, and accumulated institutional context. Many AI projects stall at the pilot phase because of this gap. Flexor's approach is to build a context layer from existing unstructured data rather than requiring organisations to manually curate knowledge bases.

By running natively within Snowflake, the integration allows organisations to process their unstructured content without it leaving their existing secure environment - a meaningful consideration for enterprises in regulated industries.

Flexor says ACE normalises, deduplicates, and standardises unstructured inputs, maps relationships between documents, emails, attachments, and conversations, and produces a consistent context layer shared across all LLMs and AI agents within a deployment.

https://flexor.ai

Agentic threats prompt IBM rethink

IBM has launched two new cybersecurity offerings aimed at helping organisations defend against a rising class of AI-generated threats, as attackers begin using frontier AI models to automate and accelerate attacks.

The company announced a new cybersecurity assessment service and IBM Autonomous Security, a multi-agent service designed to match the speed of AI-driven attacks with automated defences.

IBM warns that frontier AI models are lowering the time, cost, and expertise required to mount sophisticated attacks. Traditional security programmes built on fragmented tools and manual processes are increasingly unable to keep pace.

The new assessment service, offered through IBM Consulting, is designed to help enterprises identify security gaps, policy weaknesses, and AI-specific exposures. It delivers prioritised remediation guidance, including interim safeguards where no immediate software fix exists.

IBM Autonomous Security coordinates AI agents across an organisation's full security stack. The service analyses software exposures and runtime environments, enforces security policies, detects anomalies, and contains threats with minimal human intervention, according to the company.

IBM says the service feeds insights directly into governance and risk systems to support continuous security and compliance posture management. This claim could not be independently confirmed.

https://www.ibm.com

Lucid Links Diagramming to Claude AI

Enterprise diagramming platform Lucid Software has launched an integration with Anthropic's Claude AI, enabling users to search, create, and share Lucidchart and Lucidspark documents without leaving the Claude interface. The connector is built on the Model Context Protocol (MCP), an open standard gaining rapid adoption as a mechanism for linking AI assistants to enterprise data sources.

The Lucid Claude Connector allows Claude to locate existing diagrams, generate summaries of visual content, and convert AI-generated plans directly into editable Lucid documents.

Support for Claude Code extends the capability to software development environments, allowing developers to create architecture and process diagrams from the terminal during active coding sessions rather than after the fact.

For enterprise architects, IT managers, and digital transformation teams, the practical implication is tighter integration between AI-assisted planning and the visual documentation typically used to capture system designs, process flows, and compliance frameworks.

Lucidchart is widely used by organisations to document information architectures, data flows, and business processes - content that governance, risk, and compliance teams depend on for audit and regulatory purposes. The integration is powered by the Lucid MCP Server, which handles authentication and document access between Claude and Lucid's platform. Lucid states the connector enables documents to be shared with teammates "directly from a conversation".

https://lucid.co

Appian's answer to agentic AI risk

Enterprises seeking governance guardrails around AI automation have a new option from Appian, which has added Model Context Protocol (MCP) integration and AI-assisted spec-driven development to its process platform.

The enhancements include MCP server support for third-party AI development tools such as Claude Code and Kiro, a new Snowflake technology partnership, and a feature called Appian Composer that extracts specifications from legacy applications to guide AI-generated application development.

MCP - an open standard developed by Anthropic that allows AI agents to interact with external tools and data sources - is gaining traction among enterprise platform vendors as a means of enabling interoperability without requiring data centralisation. Appian says MCP integration will allow third-party AI agents to access its data fabric, which the company describes as providing unified read-write access to enterprise data.

The Snowflake partnership connects Appian's process orchestration layer with Snowflake's AI Data Cloud, allowing agents to query Snowflake Cortex AI directly within process workflows.

On the development tooling side, Appian Composer uses AI to extract requirements, UI models, data schemas and workflow logic from legacy systems and presents them as a visual specification. AI developer agents then work under human supervision to implement changes based on those specifications.

Appian argues this approach avoids the compliance risks and technical debt associated with unconstrained AI code generation - a concern that has surfaced repeatedly as enterprises evaluate AI-assisted development tools.

The announcement also includes enhancements to Appian's data fabric, adding a unified metadata model intended to give AI agents clearer context about how data is structured across enterprise systems. Appian has also introduced agent performance tracking and memory features that allow learnings from one process to be applied across others.

https://www.appian.com

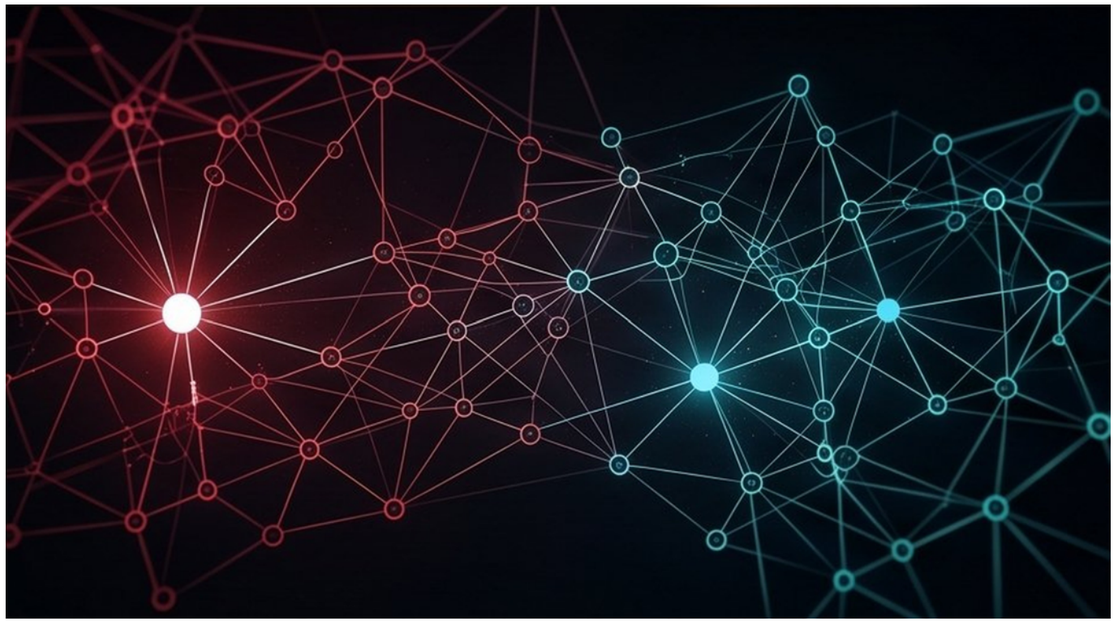

ABBYY adds Hand-written Document AI

A new alliance between document processing companies ParaScript and ABBYY brings handwriting recognition and fraud detection capabilities into ABBYY's established optical character recognition (OCR) and intelligent document processing (IDP) platform. The integration combines ABBYY's platform with ParaScript's handwriting and signature recognition, and payment fraud detection, to create a unified document intelligence workflow.

The companies say this allows organisations to process the full spectrum of document types - from structured printed forms to complex handwritten content - without replacing existing infrastructure.

"Document workflows are becoming more complex, and organizations need solutions that can adapt without increasing operational burden," said Bruce Orcutt, Chief Marketing Officer at ABBYY.

ParaScript's software development kit is already integrated with ABBYY FlexiCapture, the established on-premises IDP platform. Integration with ABBYY Vantage, the company's cloud-based IDP platform, is planned but not yet available.

The alliance targets financial services, banking, healthcare, insurance and government sectors, where high-volume document processing, fraud prevention and extraction accuracy are operationally critical.

Specific use cases include cheque and remittance processing, loan documentation, identity verification and records management requiring consistent, high-quality data extraction.

https://www.parascript.com

Sage builds AI audit trail

Sage has announced a new generation of AI agents for finance, HR, and operations workflows, led by the Sage Intacct Finance Intelligence Agent - a tool that allows users to interact with financial data using natural language and prepare tasks such as payment reminders and approvals within existing workflows.

The Finance Intelligence Agent is being released in phases and is scheduled for general availability later in 2026. It is built on Sage's financial AI models and produces a full audit trail of every AI-driven recommendation, including the data, logic, and assumptions behind each output. Sage describes this as a "glass box" approach, keeping AI recommendations transparent and explainable - a design choice driven by the regulatory and audit requirements of financial systems where, as the company's CTO Aaron Harris noted, "almost right" is not acceptable.

IDC forecasts that by 2030, 45% of organisations will orchestrate AI agents across core business functions. For finance teams, that shift requires AI that can be interrogated, overridden, and audited - not simply trusted. Sage's emphasis on explainability and logging positions the Finance Intelligence Agent for regulated environments where AI decision trails are required by compliance frameworks.

Alongside the finance agent, Sage is introducing Sales and Operational Intelligence agents for Sage X3, designed to surface risks in sales performance and operational workflows. An HCM agent for Sage's human capital management solutions targets workforce management, labour allocation, and payroll compliance.

https://www.sage.com