idm.
information & data manager
DECEMBER-JANUARY 2025

End-of-Life Security Deserves a Front Seat in Data Protection Strategies

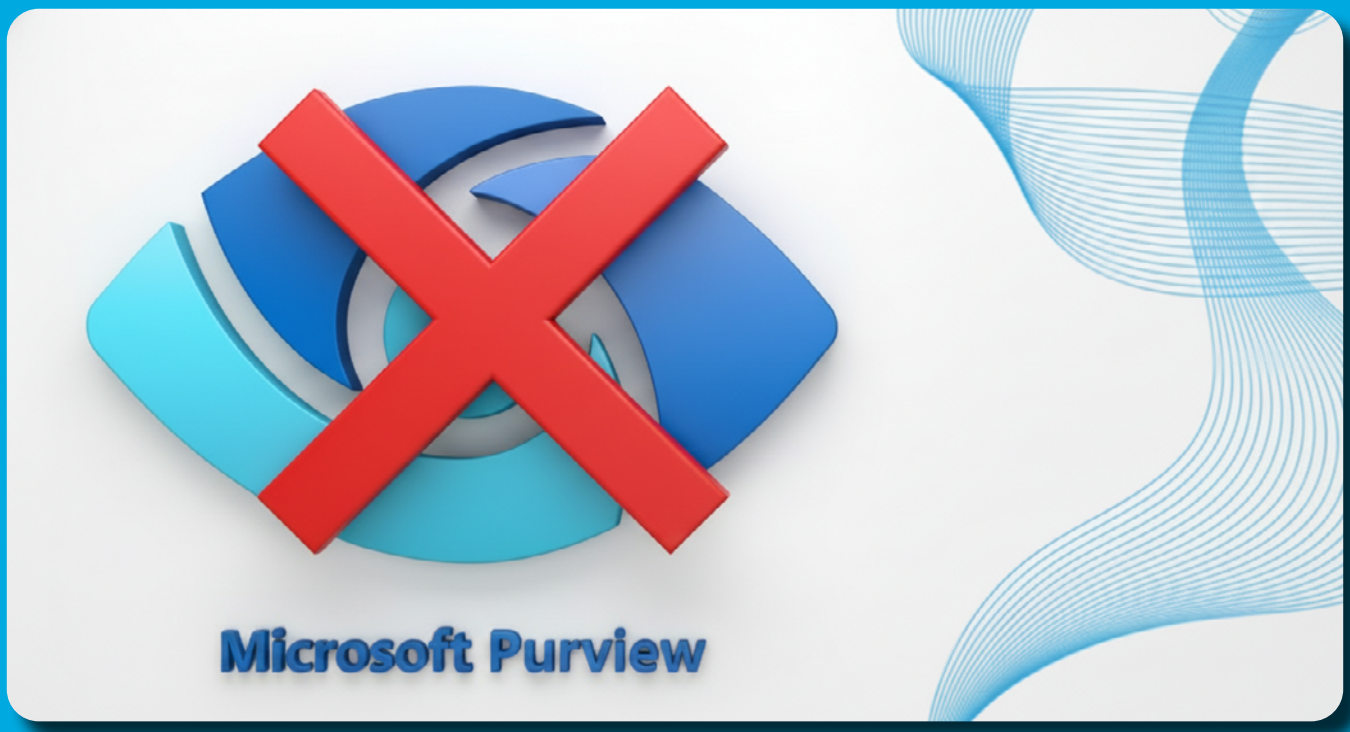

Why aren't people taking up Microsoft Purview?

Preserving Rolls-Royce Legacy Through Pixels

Government AI Survey Reveals Adoption Gaps

New Zealand Closes Privacy Law Gap

New Zealand organisations must notify individuals when their personal information is collected indirectly under new privacy legislation that takes effect May 1, 2026.

The Privacy Amendment Act introduces Information Privacy Principle 3A, requiring agencies to inform people when data is gathered from third parties rather than directly from the individual. Organisations must disclose what information was collected, the purpose, intended recipients, and individuals' rights to access and correct their data. The law addresses a regulatory gap in New Zealand's Privacy Act 2020, which previously only required notification when personal information was collected directly from individuals.

The change affects any organisation collecting data from sources such as data brokers, background check providers, referral networks, or business partners.

New Zealand's Privacy Commissioner Michael Webster said the amendment aligns the country's privacy framework with international standards in Australia, the United Kingdom and Europe, where notification for indirect collection is already required. Australia's Privacy Act requires notification under APP 5 regardless of whether information is collected directly or indirectly.

"Telling people when you're collecting information about them supports open and transparent collection practices and helps people better understand where and how their information is being used," said Mr Webster.

Organisations have until May 2026 to implement compliant systems and processes. The Office of the Privacy Commissioner will publish guidance on privacy.org.nz later this year and is reviewing existing Privacy Act codes of practice to incorporate the new requirements.

Dashboard Tool Maps Data Breaches

The Office of the Australian Information Commissioner (OAIC) has launched a new interactive dashboard that transforms how organisations access and analyse data breach statistics, revealing a 10% decrease in breaches for January-June 2025.

The Notifiable Data Breaches (NDB) statistics dashboard, unveiled today, enables users to benchmark data breach trends, assess impacts, and understand sector-specific vulnerabilities through a more dynamic interface than previous static reports. Privacy Commissioner Carly Kind said the dashboard demonstrates the OAIC's commitment to harness data for both education and enforcement purposes.

"Our goal for the new NDB dashboard is to help reporting entities learn from the experiences of others – those organisations and agencies who have had to notify us of a data breach," Commissioner Kind said.

The latest statistics show the OAIC received 532 data breach notifications in the first half of 2025, down from the previous six-month period but still indicating significant ongoing privacy risks.

Malicious or criminal attacks remained the primary source of data breaches at 59% (308 notifications), while the health sector continues to report the highest number of breaches (18%), followed by finance (14%) and Australian Government agencies (13%).

"The threat of data breaches, especially through the efforts of malicious actors, is unlikely to diminish, so we want to arm entities with data to help them keep personal information secure," Commissioner Kind said.

In her blog post accompanying the dashboard launch, Kind emphasised the importance of effective oversight when outsourcing personal information handling, stating that organisations must ensure "contractual arrangements specify accountabilities in the event of data breaches that involve multiple parties."

The dashboard builds on the OAIC's six years of experience with the NDB scheme and signals a shift toward more data-driven regulatory action, with the dashboard being updated semi-annually.

Under the NDB scheme, organisations must notify affected individuals and the OAIC when personal information has been compromised in ways likely to result in serious harm, taking reasonable steps to conduct assessments within 30 days of suspected breaches.

The OAIC has published guidance on securing personal information and data breach preparation, as well as advice for individuals responding to data breach notifications.

https://www.oaic.gov.au

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

ASIC Backs Down on Breach Data Publication

The Australian Securities and Investments Commission will not publish firm-level data on compliance breaches, reversing its April proposal following industry pushback about regulatory maturity and reporting burdens.

ASIC received 47 submissions to Consultation Paper 383 that raised concerns about publishing Reportable Situations data with company names attached. The regulator will instead publish aggregate breach data while proceeding with plans to publish firm-level Internal Dispute Resolution complaint data.

The decision reflects ongoing tensions in the RS regime, which has undergone multiple modifications since its October 2021 introduction. ASIC granted additional relief in June, extending investigation reporting timeframes from 30 to 60 days and exempting minor breaches affecting fewer than five consumers with losses under $A500.

Industry submissions warned that public naming could discourage voluntary breach reporting. The Stockbrokers and Investment Advisers Association argued ASIC would be "abrogating its responsibilities" by putting the burden of supervising licensees on consumers rather than using its investigative powers.

Compliance professionals cautioned that "name and shame" approaches could trigger under-reporting as firms avoid being publicly identified as top complaint generators.

The Australian Finance Industry Association, representing over 150 financial services firms, fundamentally opposed the proposal. "AFIA does not support publishing raw data with financial firms named as proposed under ASIC Consultation Paper 383," the organisation stated, arguing that transparency should only be introduced "where transparency can improve conduct and build trust and confidence in the financial system."

AFIA warned that without sophisticated data analytics and deeper assessment of different business models, "publication of data could be misleading, even with contextual statements." The association raised concerns that firms with strong compliance cultures that identify and report more issues could be penalised reputationally, while under-reporting firms escape scrutiny.

Law firm Herbert Smith Freehills argued ASIC has discretionary power to withhold firm-level identifiable information. The firm contended that publication "could disincentivise reporting accurate RS and IDR data to ASIC" as firms facing public disclosure may delay reporting ambiguous incidents or understate breach significance.

"There is a risk that the publication of firm-level information could increase the time taken to finalise an RS investigation," Herbert Smith Freehills submitted, noting that public data may require additional approval layers, particularly for listed companies.

The Law Council of Australia supported transparency but recommended ASIC "run a pilot program with a small number of firms of different sizes from different sectors to test the new framework" before industry-wide implementation.

ASIC's first industry-wide IDR report in October 2024 flagged data quality concerns, with 5,035 firms declaring zero complaints - higher than expected. The regulator found variations in complaint volumes among comparable firms, suggesting some may not be reporting accurately.

AFIA emphasised that "quality of data is critical for any benchmarking or comparisons exercises," noting the financial services industry contains diverse firms with different compliance systems suitable for their business models.

ASIC will publish the RS dashboard in October and the IDR dashboard later this year. The IDR publication will include firm names and Australian Financial Services Licence numbers but incorporate privacy protections for complainants and contextual explanations.

ASIC has not disclosed how many firms will be affected by publication requirements or provided guidance on data quality standards expected for public dashboards. The regulator stated it will not verify the accuracy of self-reported data

ATO Seeks AI Coding Assistant Tools

The Australian Taxation Office has launched a tender for an AI coding assistant tools to support more than 800 developers, requiring all data processing to remain onshore and comply with stringent cybersecurity frameworks. The three-year contract, with options for two additional one-year extensions, mandates the solution be hosted within Australia and prohibits offshore storage of large language models, according to tender documents.

The procurement follows growing government adoption of AI tools while addressing national security concerns about data sovereignty. The ATO specification requires compliance with Australian Signals Directorate cyber standards and the Protective Security Policy Framework.

“The solution must have guardrails that protect the ATO’s reputation by reducing the risk of harmful material being input or output from foundation models,” the tender states, requiring evidence of controls within large language model curation processes.

The procurement targets seamless integration with existing development environments including Visual Studio 2019, 2022, and Visual Studio Code, plus Azure DevOps Services Git repositories. The solution must support real-time, high-volume simultaneous usage across the ATO’s complex multi-technology application stack. Mandatory requirements include software-as-a-service delivery, encryption of data at rest and in transit, and capabilities for translating legacy code bases to modern languages. The system must also generate test scenarios and provide natural language-to-code conversion.

The tender emphasises minimal deployment complexity and comprehensive training transfer to ATO personnel. Successful providers must demonstrate adequate onshore resources and provide four-hour response times for support tickets.

Security requirements extend beyond data location to include depersonalisation capabilities for test environments and explicit guarantees that code prompts will not be retained for model training purposes. The contract value was not disclosed in available tender documentation.

UK Crown Prosecutor pilots use of GenAI

The UK's Crown Prosecution Service (CPS) has developed an AI-powered Correspondence Drafting Tool, seeking to automate complex compliance processes while maintaining regulatory oversight and data security. The CPS tool, built in partnership with NTT Data UK, utilises Microsoft OpenAI's GPT-4 models to automatically pre-populate correspondence templates with case information drawn from the organisation's existing case management system.

The CPS also looked at alternatives, such as programmatic solutions, however, due to the contextual nature of the correspondence, it was determined that Large Language Models are more suitable for correspondence drafting. Multiple LLM models were considered, however, OpenAI service was identified as the most appropriate meeting CPS security guidelines and standards.

"The tool does not make decisions in the prosecution process, it is used to aid correspondence drafting," the CPS states in its algorithmic transparency record, emphasising the critical distinction between AI assistance and automated decision-making.

The CPS said that such correspondence was previously drafted in Microsoft Word and involved manual transfers of information, which left potential for errors.

Every AI-generated draft undergoes multiple layers of human review and approval in accordance with CPS guidelines, ensuring accuracy and quality before any correspondence is sent. The organisation identified hallucination and security as primary risks, implementing specific mitigations including comprehensive user training, strict access controls, and regular quality audits.

"Data is contained within the CPS environment. The data is stored within the tool until the correspondence is finalised, then the information is deleted," according to the transparency documentation.

Currently in beta phase with 30 users, the CPS is taking a measured approach to rollout. This phased strategy has processed 230 total requests with over 440,000 tokens, providing measurable data on system performance before wider deployment.

iManage adopted by Rio Tinto Legal

Global mining company Rio Tinto has deployed iManage as part of a broader strategy to modernize and streamline its legal operations. More than 200 users across Australia, Singapore, the UK, and North America are now utilising the platform.

Seeking to improve searchability, performance, and usability, Rio Tinto replaced its legacy SharePoint-based system with iManage Work 10, Threat Manager, and Share. The implementation has delivered significant gains in user adoption and operational efficiency, while enabling the company to better govern legal knowledge and control its data environment.

The move supports Rio Tinto's focus on a connected digital ecosystem, in which integrated systems and automation simplify processes and unlock new value. This includes seamless API integration between iManage and Rio Tinto's in-house digital legal hub, with iManage serving as the core content layer.

"With iManage, we've gone from frustrated users to enthusiastic adopters," said Christopher de Waas, Digital Transformation Lead at Rio Tinto.

"People are finally able to search, file, and manage their documents without friction, and that shift has opened the door to real transformation. We're no longer just solving problems, we're building momentum."

Rio Tinto migrated over 4.5 million documents to iManage and achieved 80% user engagement within the first four months of go-live, with half of the department classified as active users. By enabling users to manage content efficiently, including while offline, iManage has helped reduce rework, increase compliance, and preserve institutional legal knowledge across the organization. Looking ahead, Rio Tinto is exploring the use of iManage's AI capabilities, particularly Ask iManage, to unlock the full potential of its legal knowledge base.

AI assist for Pacific Banking Compliance

RegGenome, a UK-based regulatory technology company, has secured a contract with the Asian Development Bank to deploy AI-powered compliance analysis across 14 Pacific Island nations. The project will compare anti-money laundering, counter-terrorism financing and cybersecurity frameworks across the Pacific region against standards in Australia, New Zealand, the United States, Japan and the EU.

Australia and New Zealand maintain robust AML/CFT frameworks under AUSTRAC and the Department of Internal Affairs respectively. Both countries serve as financial services hubs for Pacific Island nations, making regional regulatory harmonisation commercially significant. The initiative addresses declining correspondent banking relationships in the Pacific, which have restricted international financial services access. Higher compliance costs and limited regulatory information have challenged regional financial institutions.

Commenting on the engagement, Bob Wardrop, Executive Chair of RegGenome, said: "This project with ADB comes at a critical time for Pacific Island countries facing restricted access to international financial services. By providing regulators with structured data and AI-enabled tools, RegGenome is helping strengthen compliance frameworks and financial resilience while demonstrating the practical impact of computational regulation in supporting supervisors and policymakers."

The platform will deliver quarterly gap-assessment reports and training programmes for Pacific regulators. Financial institutions will be able to compare requirements across different jurisdictions more readily.

For Australian and New Zealand compliance professionals, the project offers insights into regional regulatory harmonisation and demonstrates AI applications in regulatory analysis. The outcomes may inform cross-border compliance strategies for institutions operating across the Pacific.

RegGenome is a regulatory data technology company founded at the University of Cambridge The Asian Development Bank, founded in 1966, supports development across 69 member countries.

Home Affairs Plans Manual FOI Processing Overhaul

The Department of Home Affairs has issued a request for information for a case management system to replace manual processes for the largest FOI workload in Australian government.

The department received 20,120 FOI requests in the 2024-25 financial year and assessed 2.9 million pages of documents, up from 2.1 million the previous year.

It quotes a report from the Financial Year 2021-22 when it handled 43% of all Commonwealth FOI requests. The nearest agency, Services Australia, received 14%

Current systems rely heavily on Excel spreadsheets and email-based workflows that are "inefficient, prone to error, and unsuitable for high-volume, time-sensitive work," according to tender documents.

The FOI section's existing system offers "very limited automation" while Privacy Operations manages cases entirely through Excel, a method that has proved "unreliable and unstable due to the volume of matter recording."

All document redaction is currently performed manually by officers, with no batch processing capacity available.

"The consequences of using an in-house solution which does not meet the business requirements will lead to increases in privacy breaches and to the FOI backlog," the tender states.

Privacy Operations clients are primarily internal departmental staff and Australian Border Force officers seeking support on privacy matters including complaints, data breaches and authorised disclosures to external enforcement bodies. The requested solution must provide case management, online lodgement, automated workflows, multi-format redaction with three-stage version control, and integration with existing systems including TRIM/Content Manager, ICSE, Outlook and SAP.

The system must support 130 concurrent FOI users and 20 Privacy users, handle PROTECTED classification documents, and maintain response times under two seconds during peak usage.

Required redaction capabilities must cover PDFs, emails, Microsoft Office documents, audio files, video files and image formats.

The solution must integrate with departmental repositories and provide search and retrieval across all stored artefacts and metadata.

Secure file transfer features must include time-limited links, role-based access, encryption and comprehensive audit logging to meet compliance standards under the Freedom of Information Act 1982 and Privacy Act 1988.

The tender specifies browser compatibility with Internet Explorer, Edge, Chrome, Firefox and Safari, and WCAG 2.1 Level AA accessibility compliance.

Historical data accessibility is required, with Privacy Operations needing at least three years of data and FOI Operations requiring relevant historical records for case continuity.

The two-stage procurement process began with the request for information issued 15 September 2025. Only respondents submitting RFI responses will be considered for Stage 2, the formal request for tender.

Employer Fined for Health Data Disclosure

Sydney financial services firm Fortrend Securities must pay $A13,500 in damages after deliberately disclosing a former employee's confidential medical certificate to a client. The Australian Information Commissioner found the company breached Australian Privacy Principle 6 by posting the sensitive health document to the client to discredit the departing employee.

Acting General Manager Justin Lodge ruled the disclosure was "malicious, improper and unjustifiable" in his determination. The case began when the complainant, known as 'AYN', resigned from Fortrend Securities in November 2022.

During a 30-day notice period, the complainant claimed to have experienced hostile behaviour from the company's managing director and obtained medical certificates stating they were unfit for work due to anxiety and emotional distress.

After the employee left, the managing director told multiple clients the former employee had suffered a "nervous breakdown" and was unfit to manage portfolios. When one client questioned this claim, the managing director mailed them a copy of the employee's medical certificate dated 9 December 2022.

The certificate clearly stated "PERSONAL AND IN CONFIDENCE NOT TO BE RELAYED TO ANY THIRD PARTY WITHOUT REFERENCE TO THE AUTHOR" at its top.

Lodge found the respondent collected and held the medical certificate but disclosed it for an unauthorised secondary purpose. The employee records exemption did not apply because the disclosure occurred after employment ended and served no employment-related purpose.

"The respondent disclosed the complainant's sensitive health information contained in the Medical Certificate to the Client with the intent to harm the complainant," Lodge wrote.

The determination awarded $A10,000 for non-economic loss including humiliation, hurt feelings and embarrassment. An additional $A3,500 in aggravated damages reflected the malicious nature of the breach.

Evidence from the complainant's psychiatrist confirmed they suffered ongoing anxiety and depression from the disclosure. The complainant described explaining to clients they were not having a "mental breakdown" as humiliating. Lodge noted the breach occurred within an employment relationship where the respondent held a position of trust regarding the employee's personal information. The managing director displayed "indifference towards its privacy obligations in respect of the employee's sensitive health information."

The respondent also provided unreliable information throughout the investigation, initially denying receiving the medical certificate despite email evidence showing the complainant sent it to the managing director on 12 December 2022.

Coles tackles complex compliance with automated onboarding

Coles Group has automated employment contract generation for thousands of new workers annually. The company deployed OpenText Content Management integrated with SAP SuccessFactors to replace manual, paper-based processes.

The retailer manages 130,000 employees across 1,800 stores. Each year, Coles takes on thousands of new workers. Previously, HR staff manually created and mailed contracts to candidates, causing long lead times and compliance risks when workers delayed returning signed documents.

The automated system now generates customised employment contracts using templates and business rules. Contracts are sent digitally via DocuSign and automatically filed with role-based access controls once signed.

"The combination of OpenText Content Management and SAP SuccessFactors makes our onboarding process transparent," said Bernie Repacholi, head of people and technology at Coles Group. "Candidates don't have to wait around to receive their contracts in the mail."

The implementation proved particularly valuable when new legislation took effect in August 2024. The "Closing Loopholes" reforms require employers to assess casual workers for permanent roles after specific employment periods.

For Coles, this means informing casual employees about permanent job offers within 28 days of their 12-month employment anniversary. The OpenText system automatically generates notification letters backed by full audit trails.

It also empowered Coles to dynamically customize document templates for all business units, significantly reducing manual administrative work.

The solution integrates with SAP SuccessFactors events to trigger automated workflows. When a candidate's working rights are validated through Coles' secure online portal, the system automatically pulls personal details and generates customised contracts.

When employee contracts are ready, they are sent out to candidates digitally via DocuSign. Once the candidate signs, the completed contract is automatically moved to the relevant folder in the OpenText solution, secured based on role-based access permissions from SAP SuccessFactors.

OpenText Professional Services supported the implementation and is now migrating the solution to Google Cloud Platform.

The digital process eliminates paper contracts entirely. Role-based permissions from SAP SuccessFactors control document access, strengthening governance and compliance.

Microsoft to hike enterprise SaaS cost

Microsoft will remove programmatic volume discounts for its SaaS products, including Microsoft 365, from November 1, 2025. A new Gartner report states enterprise organisations can expect list price rises of up to 13.6%. The change impacts customers with Enterprise Agreements (EAs).

The increases vary by organisation size, with the largest enterprises bearing the heaviest burden. Level D clients with 15,000 or more users face 13.6% price rises, while Level C organisations with 6,000-14,999 users see 9.9% increases.

Medium-sized Level B clients using 2,400-5,999 licences will pay 6.38% more. Level A organisations with 500-2,399 users remain unaffected, as Microsoft already removed their volume discounts in October 2018.

The changes affect Enterprise Agreements globally, with Microsoft transitioning all SaaS products to a single estimated retail price structure.

"Enterprise organisations can expect most Microsoft SaaS list prices to rise by up to 13.6%," the Gartner report states. "Unified Enterprise support contract fees will also increase, as they are calculated based on total spending."

Some products remain exempt from the changes, including US government and worldwide education price lists, software products like Windows Server and SQL Server, and Azure infrastructure services.

Gartner identifies three strategies for organisations to mitigate costs: rightsizing licence quantities through user profiling, negotiating offsetting discounts with Microsoft, and pursuing early renewal commitments before November 1.

The report recommends negotiating aggressively for discretionary discounts to offset the removal of programmatic price levels. For example, a Level D client that previously had a 15% discount might now need to secure a 25% discount off the ERP to maintain its price point. Microsoft account teams may offer one-time discounts to compensate for lost volume pricing. However, Gartner warns such additional discounts are unlikely to be repeated in future renewals.

The changes reflect Microsoft's broader strategy to provide single pricing for cloud products while potentially introducing product-specific pricing tiers in future.

For organisations with Microsoft Unified Enterprise support contracts, the price increases will compound, as support costs are calculated as a percentage of total Microsoft SaaS expenditure.

Gartner recommends preserving Enterprise Agreements despite reduced attractiveness, as they retain "From SA" price discounts for customers transitioning from on-premises software, annual licence flexibility, and superior contract terms compared to alternatives.

The research firm predicts over 60% of Microsoft's EA clients will pursue additional SaaS discounts during renewals through 2029 to compensate for lost volume pricing.

NSW Extends Metadata Contract

The NSW Government has awarded Aristotle Metadata a new three-year contract to continue hosting its central metadata platform following a competitive tender process. The platform has operated since 2018 as part of the state's whole-of-government approach to data asset management and cross-agency information sharing. The metadata registry, which supports agencies in documenting, discovering and sharing critical datasets across the public sector.

Built and managed by Aristotle Metadata, Metadata.NSW serves as a central repository where government departments can catalogue their data assets using standardised metadata protocols.

Metadata.NSW began as a proof-of-concept for sharing information between key agencies, especially around sensitive and high-impact datasets. One of the earliest successes was the Human Services Dataset, which brings together 27 years of inter-agency data about children and families in NSW. Today, Metadata.NSW is available to all NSW Government agencies, with growing uptake across departments and academic partners.

"This latest phase of the program will see us continuing on a 10-year journey with the NSW Government," said Samuel Spencer. "We're excited to welcome more agencies as part of a broad strategy to improve data sharing and data security capability across NSW."

The initiative addresses longstanding issues in government data management, where departments traditionally operated in silos with limited visibility of datasets held by other agencies.

Aristotle Metadata CEO Samuel Spencer said, "Data is more than just numbers. It's about people and outcomes, and a person-centric approach to data governance. Looking at how we can communicate across government and academia to improve citizen outcomes is critical to success."

The platform supports compliance with key legislation including the Data Sharing Act 2015 and Government Information (Public Access) Act 2009. It forms part of NSW's 2021 Data Strategy, which focuses on using data to improve citizen outcomes, treating data as a valuable asset, fostering data leadership and strengthening public trust.

The Aristotle Metadata Registry operates on the ISO/IEC 11179 international standard for metadata registries. The platform is built upon the Django web framework and allows agencies to run their own registries while maintaining federation capabilities.

The system includes the NSW Data Passport, which streamlines data access requests and supports the "5 safes" framework for secure data sharing between government and research institutions.

https://aristotlemetadata.com/

Australian Public Service Gets Blueprint for Secure AI Adoption

The Australian Public Service (APS) has been provided with a comprehensive roadmap for securely adopting Microsoft 365 Copilot, following the conclusion of a six-month government trial that demonstrated significant productivity gains from AI integration.

Technology firm AvePoint has released a detailed guide titled "*Beyond Microsoft 365 Copilot Readiness for the Australian Public Service,*" which addresses the critical security and governance challenges facing government agencies as they move from AI readiness to full-scale implementation.

The Australian Government's six-month Microsoft 365 Copilot trial, which concluded in 2024, revealed encouraging results for AI adoption in the public sector. According to the guide, 75% of users reported increased productivity, 68% saw improvements in work quality, and participants saved more than 10 work hours per month on average. However, it was identified in the whole-of-government adoption of GenAI that "Poor data security and information management practices can lead to inappropriate access to sensitive information."

Three-Pillar Security Framework

The AvePoint guide establishes a three-pillar approach for sustainable AI adoption:

■ Risk Management and Mitigation: The guide emphasises that risk management will be the biggest challenge for 82% of executives implementing AI strategies in 2025. With APS entities managing vast data volumes - some handling over 500 petabytes - the framework stresses the need for proactive threat identification and continuous monitoring.

■ Data Classification and Management: Proper data structuring and labelling are identified as essential for quality AI outputs. The guide notes that without effective classification systems, Microsoft 365 Copilot cannot deliver accurate results despite its comprehensive information access capabilities.

■ Analytics for AI Optimisation: Advanced analytics are recommended to monitor AI usage patterns, measure productivity impact, and identify successful use cases, enabling agencies to refine their AI strategies and demonstrate return on investment.

The guide emphasises alignment with established Australian security frameworks, including the Information Security Registered Assessors Program (IRAP), the Australian Cyber Security Centre's Information Security Manual (ISM), and the Essential Eight Mitigation Strategies.

AvePoint's Confidence Platform, which supports the framework, has achieved IRAP assessment to PROTECTED level, demonstrating compliance with Australian government security standards.

Addressing AI Adoption Challenges

According to Gartner research cited in the guide, Microsoft 365 Copilot introduces several security challenges including risk configuration settings enabled by default, increased oversharing risks, content sprawl risks, new retention and compliance challenges, and expanded attack surfaces requiring monitoring.

"Security professionals must simply focus on predicting steps (or missteps) that could lead to an exploit and closing those paths., said Dana Simberkoff, Chief Risk, Privacy and Information Security Officer at AvePoint

"In the absence of security education or experience, employees, users, and customers naturally make poor security decisions with technology. This means that systems need to be easy to use securely and difficult to use insecurely."

Looking ahead, the guide notes that by 2029, AI agents are expected to autonomously resolve 80% of common customer service issues, reducing operational costs by 30%. With platforms like Microsoft 365 Copilot Studio, APS entities can develop customised AI agents without complex coding requirements.

Global AI systems spending is projected to reach US$632 billion by 2028, with organisations expecting efficiency boosts (61%), increased data insights (54%), and better decision-making (51%) from AI implementation.

Rather than replacing jobs, the guide positions AI as a tool to empower APS employees to focus on strategic work requiring uniquely human skills such as empathy and ethical judgment. This represents an opportunity for workforce transformation where public servants develop new competencies in AI oversight.

AvePoint's comprehensive framework aims to help the APS transition confidently from AI readiness to effective implementation while maintaining security and compliance standards essential for government operations.

Download Guide here.

Why aren't people taking up Microsoft Purview?

By Nigel Carruthers-Taylor

Microsoft's dominance in the office applications market is undeniable - with Office 365 and Microsoft 365 (M365) powering millions of organisations worldwide, the company has established itself as the backbone of modern workplace productivity. Given this market position, one might expect Microsoft's Purview records management solution to enjoy similar widespread adoption, particularly as organisations increasingly seek integrated solutions within their existing Microsoft ecosystems.

Yet despite this natural advantage and the logical appeal of managing records within the same platform where they are created, Purview's uptake in the records management space tells a strikingly different story. Why has Microsoft's office suite supremacy failed to translate into success for its M365 records management offering, particularly in Government and highly regulated industries sectors?

It is surprising because in-place records management has been around since the mid-2000s, Purview itself has been available for nearly five years, and a number of enterprise document and records management system (EDRMS) vendors are offering significant integrations with M365 to support record-keeping.

More recently, professional associations such as RIMPA have developed education programs around Purview, and experts such as Andrew Warland have published detailed and practical guides to managing records in Microsoft 365. Yet despite these efforts, and despite the technology being available, the uptake of Purview for records management has been limited.

Surveys reinforce this sense of unease. A global study by IG World found that use of Purview for governance dropped from 50 percent to just 40 percent, with 94 percent reporting significant challenges and most using only a fraction of the available features. In some cases, the verdict was blunt: Purview, they said, is "crap for records management."

Purview's complexity and limitations

Part of the problem lies in the nature of Purview itself. Rather than being a unified, purpose-built records management application, Purview is a collection of tools - a Meccano set - that spans information governance, risk, security, and compliance. For a records manager accustomed to a structured and integrated model in an EDRMS, this can feel both overwhelming and underwhelming at the same time.

The functionality is there, but it is fragmented, technical, and often difficult to apply in a way that meets regulatory and evidentiary requirements. To manage records properly in Purview, one must first understand the complexities of M365's underlying repositories, security configurations, and authentication processes. This is not knowledge that most records managers have, nor is it easy for them to get up to speed on.

Even when organisations do engage with Purview, several weaknesses quickly become apparent. Its object-centric model applies retention at the level of individual items, such as a single email or document, rather than at the aggregate level of cases, folders, or business transactions. This undermines the ability to maintain context, which is a fundamental principle of sound record-keeping.

Worse still, when records are deleted, Purview deletes their metadata, making it impossible to prove that the record ever existed or that it was destroyed in a compliant manner. For sectors that depend on defensibility - government, finance, energy - this is not a minor flaw; it is a critical failure.

Purview's use of AI further complicates matters. While the promise of machine learning is attractive, in practice its trainable classifiers tend to apply labels at a very broad level, with little granularity.

This makes it difficult to know whether the correct retention schedule has been applied to any given record and can create compliance risks if high-value records are inadvertently treated as low-value. Instead of providing assurance, this broad-brush approach can undermine confidence in the system and force records managers to double-check the outcomes manually - negating the efficiency gains that AI was meant to deliver.

The trouble with "in-place"

These weaknesses originate from Purview's use of the manage-in-place model. In theory, letting records remain where they are created - in Teams, Outlook, or OneDrive - minimises user burden and handles scale. In practice, however, it creates fragmentation and duplication.

Academic analysis concludes that "the in-place model involves acceptance of sub-optimal structure/schemas, so in circumstances where it is possible to optimise the efficiency of a structure/schema of a corporate records system we should reject that model." *(Lappin, Jackson, Matthews, et al).* This rejection stems from the obvious risk of a single record, or versions of it, being duplicated across multiple locations with inconsistent security, which undermines authenticity and drives up the cost of discovery, FOI requests, and audits.

More importantly, scattered records lack context and fracture the evidential trail of business transactions, leaving gaps that force AI classifiers to "make up stories" to compensate. Given that Microsoft itself has promoted the use of AI for classification in Purview, this raises real questions about reliability and trust.

It is no wonder, then, that many organisations hesitate to adopt Purview as their sole records management solution. Professionals understand that while Purview has strengths in governance at scale and compliance monitoring, it lacks the lifecycle controls, contextual integrity, and defensible disposal practices required for full records management.

This is why traditional EDRMS' remain trusted, especially in government - they maintain records in context, are inclusive of hardcopy, preserve metadata, enforce classification, and ensure evidentiary defensibility through hybrid models. Nonetheless, Purview's ability to auto-classify and apply retention at scale shouldn't be dismissed, as it helps tackle the information overload facing most organisations. So, what's best practice?

Where hybrid strategies come in

It is here, in this space of complexity, that EDRMS' demonstrate their continuing value. These platforms build on decades of proven capability and most now offer to extend that relevance into the modern world of M365. These new EDRMS capabilities can hide and automate much of the mechanics of record-keeping, enhance manage-in-place, simplify the user experience, and integrate with M365 and other business systems.

Critically, they support a hybrid model that draws together three established traditions of records management: the centralised model associated with Duranti (advocating central control to preserve authenticity), the integrated model articulated by Bearman (embedding record-keeping within business systems), and the in-place model familiar to most M365 users. By combining these approaches, an EDRMS provides a way to take advantage of Purview where it works, while also addressing its shortcomings.

The practical implementation of this hybrid model is relatively straightforward. Low-value or short-term records, where risk is minimal, can remain governed in place through Purview. High-value records, however, may be identified by Purview using retention labels - whether applied manually or automatically through machine learning.

Once identified, these records can be uplifted to the EDRMS for stronger lifecycle management. In some cases, the record may still be "managed in place" but with the EDRMS applying robust classification, retention, and disposal controls. In others, it may be managed or moved into the EDRMS to sit within the appropriate business context.

In the Australian Government context, OpenText Content Manager remains the most widely recognised and compliant EDRMS, providing the authoritative control for high-value records. Solutions like Ingress by iCognition extend this further by enabling seamless synchronisation between Content Manager and Microsoft 365, allowing Purview's labelling to trigger Content Manager's more rigorous compliance and disposal processes.

In this way, organisations can balance the scalability of in-place governance with the evidentiary assurance of a centralised EDRMS, overcoming Purview's limitations in three critical areas: the lack of granular classification, the absence of defensible disposal records, and the risks inherent in scattered, duplicate repositories.

A balanced way forward

This hybrid model has significant implications. It means that Purview can be used for what it does well - broad in-place governance at scale across SharePoint, Exchange, and Teams - while EDRMS', such as Content Manager via Ingress, ensures that high-value records are subject to defensible, standards-compliant lifecycle management.

Rather than duplicating everything or overburdening users, the system manages records according to their value, preservation and compliance requirements. In doing so, it balances risk, reduces cost, and aligns with best practice in Australian Government standards and international frameworks alike.

Importantly, Purview is not the only in-place management system that faces in-place issues and limitations. Other platforms, such as Castlepoint or OpenText CDDRI, also operate in-place and deliver real value, but like Purview, they achieve best practice when complemented by an EDRMS such as Content Manager to provide evidentiary assurance, structured metadata, and compliant disposal.

The question of why Purview has not been widely adopted, therefore, may have a simple answer: it is not fit on its own to carry the weight of modern records management. But that does not mean it should be dismissed. Instead, it should be used as one part of a hybrid solution that integrates with proven systems and best practice models.

Purview identifies and labels; EDRMS' like Ingress and Content Manager classify, control, and dispose. Together they create a model that respects the realities of modern digital environments while meeting the enduring demands of compliance and accountability.

In conclusion, Purview's limited uptake is not simply a matter of education or promotion. It reflects deeper concerns about its complexity, architecture and compliance adequacy. By recognising these limitations and adopting a hybrid model, organisations can move beyond the false choice of "in-place or centralised" and instead embrace a balanced approach. In the era of Microsoft 365, hybrid records management is not just a strategy - it is the practical way forward.

Nigel Carruthers-Taylor is Executive Director & Principal at Information Management and Governance Specialists, iCognition. For more information contact iCognition on info@icognition.com.au.

EzeScan Launches Automated PII and PCI Discovery Suite

Organisations now have a scalable solution to locate and secure sensitive peronal data buried across sprawling digital repositories. EzeScan's Document Repository Analyser (DRA) – part of its PII & PCI Automated Discovery & Redaction suite – delivers both discovery and remediation capabilities integrated in a single platform.

The solution addresses a critical gap in compliance workflows. Most discovery products on the market can only search and flag sensitive data – they cannot automatically remediate it. Organisations face mounting pressure to find personally identifiable information and payment card data whilst staying compliant with tightening privacy legislation.

"What our solution does that the other solutions on the market don't do, is we can search a repository, we can find the PII data and then we can align with a compliance workflow to remediate that," said Demos Gougoulas, Director of Sales and Marketing at EzeScan.

"Our competitors would force their users to look at every single file and manually redact it."

EzeScan's DRA interrogates multiple enterprise systems critical to compliance teams: Content Manager, Objective, SharePoint, OneDrive and network file shares. The platform discovers duplicates, detects classification errors and identifies sensitive data across these repositories simultaneously.

Demos Gougoulas, Director of Sales and Marketing at EzeScan. EzeScan has been in business since 2002, with thousands of installations across Australasia, North America and the United Kingdom, providing capture and business process automation solutions including robotic process automation, forms data extraction and integrated EDRMS capture.

"We can do it in situ," Gougoulas explained. "We can scan your Content Manager repository or your SharePoint libraries or your network drives with software that's installed on-premise to go and look for it. When it finds it, it provides you a workflow to remediate it."

This on-site capability addresses a significant barrier to compliance adoption. Most competitor products operate cloud-based models requiring organisations to upload documents for analysis – an impractical scenario for agencies managing millions of documents.

The solution offers deployment flexibility critical to large organisations managing sensitive information. "We offer the solution on-premise or in the cloud and it can act on ECM or file storage on-premise or in the cloud," Gougoulas noted.

"We can do cloud-to-cloud if we have to. Some of the AI engines are external. If we need to be looking for handwriting, we use Microsoft's handwriting recognition engine, and then the customer can consume that themselves within their own Azure environment."

This approach lets Governance Risk and Compliance Managers maintain data sovereignty whilst leveraging advanced detection capabilities.

EzeScan employs language models and image-based AI engines to locate sensitive information. The solution identifies images of passports, driver's licences and Medicare cards – even when photocopied at angles or partially obscured.

For organisations with legacy document repositories, this capability proves critical. One state government entity managing two million documents knew photocopied passports existed within their collections but lacked tools to locate them. EzeScan's PII/PCI compliance solution discovers what traditional OCR-dependent systems cannot.

FOI and Email Compliance

The platform streamlines Freedom of Information (FOI) workflows – a priority for government agencies and organisations subject to information access legislation. Records Managers and FOI officers can create FOI workflows within EzeScan's web application, enabling document approval, temporary redactions and audit trails.

Sensitive data protection extends to email systems where compliance risks frequently hide. For example, if an organisation is at risk of its staff forwarding email with sensitive information, EzeScan can find those emails and make the user aware of them before they push the send button.

"Our software will check the emails, and it will put a tag on the e-mail that says PCI or PII detected," Gougoulas explained. "So, there's a visual cue for the user to go in and delete it. If required, we can convert an e-mail to a PDF and then redact it, and then we can save it as a record in the EDRMS."

The platform's remediation capabilities extend beyond simple redaction. When credit card data appears on network drives, EzeScan's workflow can automatically redact the sensitive content, move the redacted version to an EDRMS like Content Manager, and place a stub on the original network location referencing the controlled copy's new location.

If business requirements mandate retaining original documents with sensitive data, the system can securely archive those separately whilst maintaining compliance documentation. Records Managers can bulk-search and redact Freedom of Information requests containing thousands of pages, applying reason codes for audit compliance.

Addressing Budget-Constrained Compliance Programs

Competitive pricing has emerged as a significant differentiator. Government agencies report that existing solutions on the market command "ridiculous sums of money," leaving many organisations without budgets to address compliance risks until security incidents force action.

"Our aim is to make this tool available to everybody without having to sell the farm," said Gougoulas.

EzeScan's subscription model and on-premise installation approach keeps costs down and helps organisations with their procurement.

EzeScan's solution builds on its legacy in document capture, scanning and image-based processing. The company has spent years developing PII and PCI detection capabilities following requests from government entities facing compliance challenges after not finding what they needed in the market.

"They are ecstatic because to buy a similar product, which doesn't remediate, would cost them a couple hundred thousand dollars," Gougoulas said of early deployments. "Ours is nowhere near that."

"It also provides a lot of added functionality, so while its searching for PII it can find and flag duplicate documents, OCR any documents that are not text searchable, or other business orientated workflows.

"So, it's a lot more feature rich to suit information management requirements."

https://www.ezescan.com.au/solutions/pii-and-pci-automated-discovery-and-redaction

TRIM/CM Hits 40-year Milestone

OpenText has unveiled an AI-powered roadmap for its Content Manager platform, which celebrates 40 years of securing government information assets across Australia and globally in 2025.

Originally developed in Canberra in 1985 as TRIM (Tower Records Information Management), the platform now supports more than 700 organisations and 1.8 million users worldwide, from the National Archives to local councils.

OpenText has unveiled its 2025-2026 roadmap for OpenText Content Manager, focusing on modernisation, intelligent automation, and seamless integration with platforms like Microsoft Teams, SharePoint, and Google Drive.

Key updates include a modern interface, enhanced mobile app features, AI-driven natural language processing for simplified searches, advanced reporting tools, and improvements to auto-classification, empowering organisations with smarter content management and actionable insights.

OpenText ANZ Vice President George Harb said what began as a records management system for public servants has evolved into a modern, cloud-ready information governance platform, maintaining its reputation for security, compliance and trust.

"Content Manager's evolution mirrors the growth of digital government itself, from paper archives to AI-driven information ecosystems. It's a testament to how innovation, compliance and trust can coexist over decades," Mr Harb said.

"As agencies face growing pressure around cybersecurity, privacy reform and transparency, Content Manager continues to give them confidence that critical records are secure and discoverable.

The platform remains deeply embedded in Australia's public sector while expanding internationally to organisations like the U.S. Department of Energy's Office of Legacy Management, which uses Content Manager for long-term preservation of environmental and regulatory data.

According to Brandon Voight, OpenText Australia and New Zealand Director Public Sector, the platform's longevity stems from partnerships across Australia and New Zealand that have continuously enhanced its capabilities.

The software was originally built by Tower Software, founded in 1986 by Brand Hoff, with Rory Kleeman heading research and development. Kleeman credits the platform's success to addressing the fundamental need for information accountability and governance.

"From the start, it was about giving people confidence that their information was safe, accessible and governed properly. That's what made it stick - not just the technology, but the mindset," Kleeman said.

FOI Reform Ignores AI Implementation Issues

Australia's proposed Freedom of Information Amendment Bill 2025 fails to address how artificial intelligence should be used in making or processing government transparency requests, according to a new analysis by law firm King & Wood Mallesons that identifies critical gaps in the legislation.

The Bill, introduced into the House of Representatives last week, aims to "modernise the requirements for Freedom of Information requests" but remains silent on AI's role in FOI processes, write senior associate Kendall Mutton and partner Rebekha Pattison in their legal analysis.

"Curiously, the Bill is silent on the use of AI systems to make or process freedom of information requests," the authors note, despite AI already being implemented across multiple Australian Government agencies.

The lawyers identify three key areas where the legislation creates uncertainty: whether AI can autonomously make FOI requests, how to process requests involving AI-generated government decisions, and using AI to handle FOI applications.

"It is conceivable that an AI system could be trained to itself make an FOI request," Mutton and Pattison write, but conclude that "the better view for the current compilation of the FOI Act is that AI cannot itself make a valid FOI request as AI is not a 'person'."

A significant compliance question emerges around whether AI-generated government information constitutes a "document" subject to FOI disclosure. Under the FOI Act, documents are broadly defined to include "any article on which information has been stored or recorded, either mechanically or electronically, as well as any other record of information."

Crucially, the lawyers note there is "no requirement in the FOI Act for a document to have been created by a 'person'" and this gap isn't addressed in the Bill. They conclude that "AI-generated information seems likely to be considered an article on which information has been stored or recorded, either mechanically or electronically, for the purpose of the FOI Act."

This interpretation gains support from case law in other jurisdictions. In DPP v Khan, the Supreme Court of the Australian Capital Territory referred to AI-generated character references as "documents," though "there was no detailed commentary on whether and why AI-generated material constitutes a document."

However, "whether the underlying algorithm used to generate the information will be similarly caught is a trickier question," the analysis states. This depends on algorithm content and request scope – for example, if an FOI request seeks documents containing specific terms, it "could conceivably include an AI algorithm which uses that term, if the algorithm is considered to be a document."

The analysis highlights practical enforcement challenges, noting: "As a practical matter, unless the FOI Act is amended or an agency requests that the use of AI is disclosed when making a request, it may be difficult to identify when a request has been generated using AI."

For government agencies implementing AI systems, the legal experts identify "one exciting opportunity" in using AI to generate "significant efficiencies" in processing FOI requests. United States government agencies have been testing AI since mid-2023 to perform keyword searches, summarise document characteristics, and identify potential exemptions.

However, the Bill doesn't clarify whether automated systems can make FOI decisions. "It is not clear whether this was a deliberate decision to keep this type of administrative action in the hands of human decision makers given the levels of judgment that are often required, or whether this was a missed opportunity to open the door for more efficient processing," the authors observe.

View the original article here

$A83M contract backs NSW digital patient record rollout

The NSW Single Digital Patient Record Implementation Authority (SDPRIA) has awarded a contract worth $A83 million to RLDatix Galen Australia for a statewide data archive solution.

The contract, effective from 27 December 2024, will support the migration and management of historical health data as NSW Health implements its ambitious Electronic Medical Record (eMR) modernisation program. The arrangement extends to 16 December 2034.

RLDatix Galen's data archive solution forms a critical component of the broader Single Digital Patient Record program.

The platform will store and manage data from NSW Health's existing systems as the state transitions to an integrated Epic Systems-based electronic medical record across all public hospitals, community health centres, and pathology laboratories.

The program's scale is substantial. The Single Digital Patient Record will eventually serve 228 public hospitals, more than 600 community health centres, 60 pathology laboratories, and over 150 pathology collection centres. It will replace multiple existing systems used across all 17 local health districts and specialty health networks.

NSW Health is not developing the technology independently. Epic Systems provides the core eMR platform, while Amazon Web Services supplies the secure cloud hosting environment.

RLDatix Galen's archive solution completes the core infrastructure partnership for data management.

NSW's SDPR initiative occurs within a broader, fragmented Australian healthcare IT landscape where states have adopted diverse vendor strategies and implementation timelines.

South Australia has standardised on Allscripts Sunrise EMR and PAS, with the government approving additional funding for a statewide regional rollout expected to be complete by late 2024.

Tasmania has completed procurement and selected Epic Systems as the preferred EMR vendor, with contract negotiations underway and hopes to begin building a statewide Epic EMR solution in May 2026.

Victoria has pursued a more fragmented approach, with Altera Digital Health receiving positive feedback for its cloud-based EHR solutions in regional and remote healthcare settings, particularly in Gippsland where telehealth services grew by 40% from 2022 to 2024.

The Victorian government is investing A$21.4 million to support four health services to transition to electronic records, including the Royal Eye and Ear Hospital, Eastern Health, the Hume Rural Health Alliance, and Grampians Rural Health Service.

Queensland has increasingly adopted Telstra Health's Kyra Clinical solution, which facilitates realtime data sharing across hospital networks and virtual care systems, with Telstra Health recognised for its focus on interoperability and early adoption of FHIR (Fast Healthcare Interoperability Resources) standards.

New Zealand presents a contrasting model to Australian state standardisation efforts.

Rather than pursuing single vendor consolidation, Health New Zealand has embarked on the Shared Digital Health Records (SDHR) project to connect data from existing shared digital health records and nationally available clinical data into a consistent view, leveraging existing access, consent, and privacy controls, initially funded with NZ$4 million through its launch in 2025.

New Zealand's Ministry of Health has moved away from building a single Electronic Health Record towards developing a national Health Information Platform that will enable data about a single patient to be shared, with the Ministry planning a phased approach to implementation with investment in tranches and avoiding 'lock in' to a single technology solution.

My Health Record Integration

A critical dimension of Australian state EHR deployments involves integration with the national My Health Record system (previously called PCEHR). State-level EMR implementations are increasingly required to support bidirectional data exchange with the national platform.

The Australian Digital Health Agency completed the final stages of integration between South Australia Health's Sunrise EMR and patient administration system to the country's My Health Record, with an embedded tab within the Sunrise EMR providing clinicians with access to MHR which creates a unified view of a patient's interactions across the health care system containing shared health summaries from general practitioners, pathology and imaging reports as well as prescription information from a patient's visit both within South Australia and interstate.

HealtheNet, NSW Health's information sharing platform and secure clinical portal, receives and shares clinical information across NSW Health facilities and My Health Record, and provides NSW Health clinicians with immediate access to an aggregated view of their patient's health information, including information which resides outside of the public hospital system.

When NSW Health patients visit hospitals or health services, discharge summaries, pathology test results, discharge dispense medication and diagnostic imaging reports will be sent to their My Health Record unless they choose not to have this information sent.

Telstra Health's Kyra Clinical natively supports uploading discharge summaries to My Health Record to meet Australian Digital Health Agency requirements, with the system integrating radiology and pathology systems for electronic ordering and viewing of results.

Compliance requirements around My Health Record uploads vary by jurisdiction. In NSW, patient consent is required to upload specific health information to My Health Record under state legislation, as is the case in ACT and Queensland per the My Health Records Regulation 2012. This creates operational complexity for health information managers implementing systems across multiple states.

SDPR Rollout Timeline

Implementation of the NSW Single Digital Patient Record will proceed in five tranches, with Tranche A launching in March 2026 across Hunter New England Local Health District, NSW Health Pathology, and Justice Health & Forensic Mental Health Network.

Subsequent tranches will progressively roll out throughout 2026, 2027 and mid-2028, with Tranche B (late 2026) involving Northern NSW LHD, Mid North Coast LHD, Northern Sydney LHD, Central Coast LHD, and LIMS North, followed by further tranches covering remaining local health districts.

Both Hunter New England LHD and Justice and Forensic Mental Health Network currently use systems from Orion Health, which will be replaced by the Epic-based SDPR. This represents a significant operational transition for the 700+ bed regional network.

NSW Health is currently undertaking ongoing readiness activities and testing. Training programs for staff will intensify closer to each stage of rollout, though specific implementation timelines by individual health districts have not been publicly detailed.

Details about the volume of historical data to be migrated from Orion Health, expected system availability windows during transition, or service level agreements for archive platform uptime during the March 2026 launch remain opaque from procurement disclosures.

AI Procurement Guide for Australian Organisations

Australian law firm MinterEllison has released comprehensive guidance for organisations procuring artificial intelligence systems, emphasising risk-based classification and robust contractual protections as businesses increasingly integrate AI into operations.

The guide addresses critical procurement considerations including data ownership, liability management, and compliance with Australia's strengthened privacy laws, which now require organisations to disclose automated decision-making processes in privacy policies.

"AI procurement must be risk-classified to determine due diligence, governance, and contract protections, especially for high-risk use cases like legal or personal data processing," the guidance states.

The framework categorises AI systems by organisational risk levels - high, low, or exempt - based on intended use, data sensitivity, and potential impact on individuals or compliance obligations. High-risk systems require independent audits, clear training data documentation, and ongoing monitoring mechanisms.

The guidance outlines essential contractual clauses organisations should negotiate, including:

- Clear data ownership definitions preventing AI systems from training on organisational information
- Indemnities for legal breaches, intellectual property infringement, and AI-generated harm
- Transparency requirements including model cards and decision logic summaries
- Human oversight provisions allowing review and override of automated decisions
- Security standards and incident response protocols

Australian Privacy Law Implications

The guide highlights increased risks under Australia's recent privacy law reforms. These changes require organisations to update privacy policies when using automated decision-making processes, reflecting growing regulatory concern over AI and personal data use.

Organisations face potential fines or litigation if privacy laws are violated, particularly when AI vendors seek to use client data for system improvements that could expose information elsewhere.

The guidance warns that AI's use of large datasets has significantly increased intellectual property infringement risks. This occurs when copyrighted data enters AI systems or when outputs reproduce copyrighted training material.

While organisations can typically shift infringement risks to vendors through indemnities, resistant AI vendors may require more extensive due diligence into training datasets.

The procurement framework emphasises thorough vendor assessment, including verification of internal governance frameworks, audit trails, and accountability mechanisms. Organisations should seek independent certifications such as ISO/IEC 42001, signalling commitment to responsible AI development.

The guide recommends confirming vendors can provide clear documentation about AI system functionality, limitations, training frequency, and data sources - whether live internet data or contained organisational datasets.

Implementation Checklist

MinterEllison also provides a practical checklist covering eight key contractual areas: data use and ownership, liability and indemnity, performance standards, transparency requirements, regulatory compliance, human oversight, security protocols, and termination rights.

The guidance acknowledges organisations' negotiating position with AI vendors will vary but emphasises incorporating these considerations early in procurement processes to safeguard data and ensure compliance.

This guidance addresses growing demand for structured AI procurement approaches as Australian organisations balance digital transformation opportunities with complex legal and ethical obligations.

View the guide here.

BCA Calls for Medical Records Digitalisation

Australia's healthcare system requires urgent digital transformation to address mounting cost pressures and workforce shortages, with technology investments potentially saving $A5.4 billion annually, according to a new Business Council of Australia report.

The 168-page blueprint calls for mandatory interoperability standards across health, aged care and disability sectors by 2028, warning that fragmented digital systems are hampering productivity and patient outcomes.

It has also suggested the introduction of a Digital Health and Care Interoperability Incentive Fund within the next 3 years to support the digitisation of paper-based records and interoperability across the health and care system.

“In 2025, it would astound many Australians that the majority of hospitals are still paper-based and that up to 75 per cent of global fax traffic comes from medical services,” said Rohan Mead, chair of the BCA’s Health and Care Services Committee.

The report identifies critical gaps in Australia’s health technology infrastructure, noting that 90 per cent of OECD countries have electronic health portals, but only 42 per cent allow public access to all their data.

The BCA proposes renaming the Australian Digital Health Agency to the Australian Digital Health and Care Agency, expanding its remit to oversee digital transformation across aged care and disability services.

Key recommendations include establishing a public register of health software that integrates with My Health Record, creating a Digital Health and Care Interoperability Fund, and mandating the National Healthcare Interoperability Plan 2023-2028 across all sectors.

The report estimates that electronic medical records across public hospitals could save $A355 million annually by reducing duplicated pathology and imaging tests.

Artificial intelligence could automate up to 30 per cent of healthcare tasks, potentially freeing clinicians for patient care, the report states. However, it warns that regulatory frameworks must keep pace with technological advancement.

“AI has the potential to free up 30 per cent of a clinician’s time, allowing them to spend more time with Australians,” Mead said.

The report projects a shortage of 79,000 nurses by 2035, with digital solutions and scope-of-practice reforms identified as partial solutions to workforce pressures.

Current fragmentation sees separate portals for My Health Record, aged care, NDIS and carer services, creating data silos and administrative inefficiencies.

The blueprint calls for a coordinated national approach, proposing establishment of an Australian Health and Care Commission to consolidate existing regulatory bodies and harmonise standards.

Government health and care expenditure increased from $A70 billion to $A184 billion between 2011-12 and 2023-24, with the sector projected to exceed 10 per cent of GDP by mid-century.

Governance Issues Top Barrier to AI Success

Poor governance and data management practices, not technological limitations, are the primary reasons artificial intelligence initiatives struggle in organisations, according to new research that challenges recent claims about widespread AI project failures.

A survey of more than 200 business and IT leaders by The Data Warehousing Institute (TDWI) found governance issues were cited by 49% of respondents as their greatest AI frustration. This was followed closely by AI hallucinations at 46% and lack of AI literacy at 48%.

The research directly counters headlines from a recent MIT study claiming 95% of AI projects fail. TDWI researchers argue their findings complete the narrative by highlighting organisational rather than purely technical barriers.

"The roadblocks to AI success aren't just technical; they're organisational," said Meighan Berberich, TDWI president. "Companies that put governance, literacy, and a builder mindset at the centre of their strategy won't just survive the AI transition, they'll thrive."

The study distinguished between "consumers" who use packaged AI tools and "builders" who embed AI into workflows with their own data. Builders reported faster decision-making (64%) and increased innovation (46%), while consumers mainly cited time savings.

Despite implementation challenges, 90% of survey respondents already use general-purpose generative AI assistants such as ChatGPT and Perplexity. Nearly half of builder organisations are experimenting with workflow automation including onboarding systems and incident reporting.

“Too many organizations are just ‘paving the cow path’ with AI, applying a thin layer of automation over broken processes and expecting transformative results,” said Tamilla Triantoro, PhD, professor of business analytics at Quinnipiac University and TDWI contributor.

“If you aren’t prepared to redesign your workflows, govern your data, and upskill your people, you aren’t driving transformation - you’re accelerating inefficiency.”

The findings emphasise that organisations must address data governance, skills development and process redesign before expecting measurable returns on AI investments. This aligns with regulatory frameworks emerging globally around AI governance and data management.

The research brief is available at www.tdwi.org. The survey was conducted in June and July 2025 with 155 qualifying responses from various company sizes and industries.

The Perils of Document Fraud

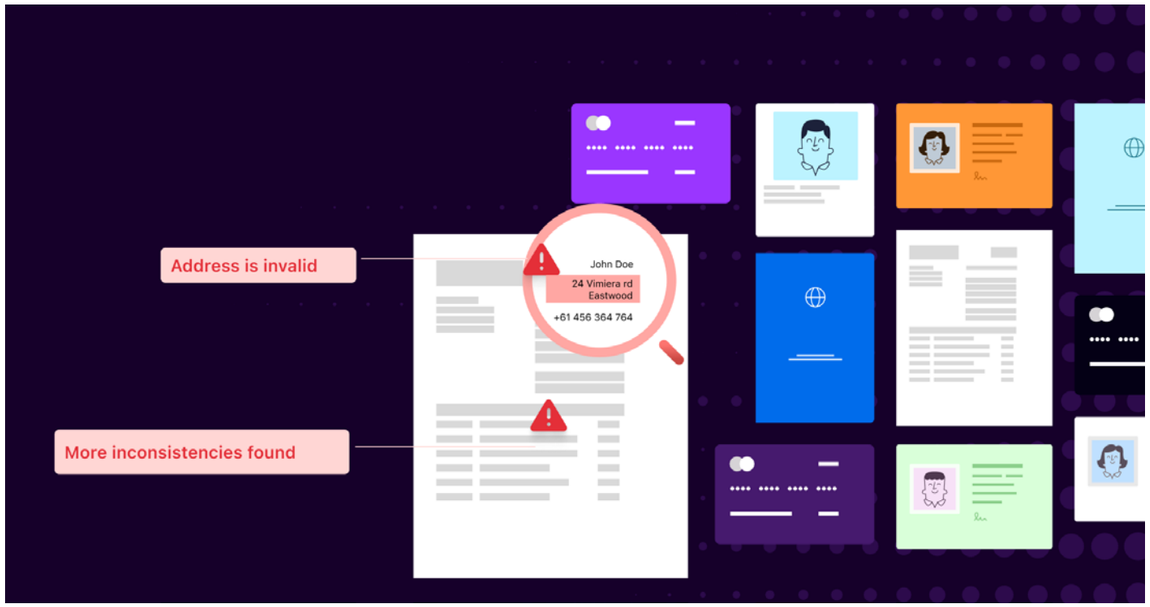

In October 2025, Noosa Council CEO Larry Sengstock revealed a sobering truth: the council had fallen victim to a sophisticated $A1.9 million fraud perpetrated by international criminals. Despite having standard processes in place, the organisation was targeted through what Sengstock described as "social engineering AI techniques" – a sophisticated, strategic attack that bypassed existing safeguards.

"We can reveal that the fraudulent activity was sophisticated, strategic, and targeted," Sengstock stated. "Unfortunately, in this instance [our processes] were not effective enough, as this crime was committed by highly organised, professional criminals who found a way through our processes."

This incident isn't isolated. According to the Australian Bureau of Statistics' 2025 Personal Fraud Survey, identity theft affected approximately 255,100 Australians in the 2023-24 financial year. While CEO Sengstock did not explicitly detail the nature of the fraud due to ongoing investigations, document fraud was likely a central component in the attackers' strategy. Financial fraud of this magnitude typically involves falsified payment authorisations, manipulated invoices, or forged vendor documentation that bypasses standard verification procedures.

As Paco Romero Ferrero, a data scientist specialising in verification automation at Veriff, explains: "Subtle indicators, such as mismatched fonts or incorrect borders, can signal a fraudulent document." Ferrero emphasises "the crucial role of collaboration between Fraud Operations and Engineering Teams in staying ahead of evolving fraud tactics" by monitoring trends and continuously updating detection capabilities.

"This pattern of sophisticated document manipulation aligns with the growing trend of document-based attacks on government entities, where criminals exploit gaps in verification processes rather than attempting more technically challenging system breaches.

AI-Powered Document Fraud: A Growing Business Risk

Document fraud has evolved dramatically in recent years. The Entrust 2025 Identity Fraud Report reveals that digital document forgeries increased by a staggering 244% last year alone. As the Australian Federal Police warned in October 2025, more than 40% of all Australian cybercrime victims fall prey to multiple types of cybercrime in a single year, with fraud victims emerging as the most vulnerable to repeat targeting.

"Today's cybercriminals can download document templates or breach-obtained images, then alter them using common tools like Photoshop or even generative AI to quickly create synthetic identities," notes cybersecurity expert Mason Wilder in the Association of Certified Fraud Examiners (ACFE) 's "Top 5 Fraud Trends of 2025" report.

"In 2025, AI tools will be applied to the creation and use of synthetic identities in fraud schemes targeting government and private organisations at a higher rate."

Introducing DoxAI's Fraud Check AI

In this increasingly threatening landscape, DoxAI's Fraud Check AI offers Australian organisations a powerful defence against document fraud.

DoxAI is an artificial intelligence-focused company headquartered in Sydney, focused on serving industries where regulatory compliance, document verification, and secure data handling are critical concerns.

DoxAI's comprehensive Fraud Check AI solution performs intelligent fraud checks by analysing document structure, metadata, and digital fingerprints to identify manipulation or forgery, delivering accurate results through advanced AI risk scoring.

Unlike traditional verification methods, Fraud Check AI:

- Detects AI-generated documents instantly
- Performs comprehensive metadata analysis to identify origin, authorship, and modification history
- Validates handwritten and electronic signatures
- Cross-checks information against trusted external databases
- Analyses raw file structure for unauthorised modifications
- Applies category-specific verification rules for different document types

The solution supports over 170 document categories out-of-the-box, including financial statements, payslips, tax documents, invoices, and government-issued identification – providing comprehensive protection across all departments.

Measurable Business Benefits

- Beyond security, Fraud Check AI delivers tangible operational advantages:
- **120x faster** than manual fraud analysis, completing verifications in seconds instead of hours
- **50x more cost-effective** than manual fraud checking, saving on labour and operational expenses
- **-87% fewer errors** than human fraud analysis, drastically reducing costly mistakes

With SOC2 Type 2, PCI DSS, GDPR, ISO27001 compliance and enterprise-grade security features, DoxAI's solution ensures both robust protection and regulatory compliance.

As Noosa Council and countless other organisations have learned, traditional verification processes are increasingly vulnerable to sophisticated fraud attempts. Don't wait until after an incident to strengthen your defences. Join us to learn how AI-powered document verification can protect your organisation from today's most advanced fraud techniques.

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

Startup Addresses Messaging Risks

Enterprise leaders face a critical compliance challenge as workplace communication fragments across personal and corporate platforms, creating data security and operational control problems that existing solutions fail to address. An Australian–New Zealand startup plans to fill this gap with locally controlled infrastructure.

Corro, launched by Xero founder Rod Drury and former AWS executive Sara Goldsworthy, is developing a secure messaging platform designed to give organisations control over sensitive workplace communications. The platform prioritises data sovereignty alongside security, directly addressing the trust gaps many compliance–focused organisations face when using consumer–grade communication tools.

"Communication is the connective tissue of society, the invisible infrastructure of organisations," Goldsworthy stated. "Yet we're forced to work around our systems rather than them working for us."

Research shows 67% of employees use personal communication applications at work, while 60% report difficulty locating messages they know exist. These behaviours create compliance exposure for organisations managing sensitive information in regulated industries including government, finance, healthcare, and legal services.

A Proofpoint report published in November 2025 found organisations experience an average of 11 data loss incidents yearly, with nearly 58% attributable to careless employees or contractors. Enterprise chat remains a significant blind spot: 87% of enterprise chat usage occurs through unmanaged accounts, and 62% of users paste personally identifiable information into them.

"Leaders are trapped in an impossible trade–off," Goldsworthy explained, highlighting the tension between security and usability. "Lock everything down and lose flexibility, or allow choice and live with chaos. We're not accepting that as inevitable. The stakes go beyond productivity—it's about protecting institutional knowledge, maintaining operational security, and ensuring control over critical communications."

Corro is being built by a team that combines government security expertise with private–sector technology experience. Head of Product Jess Modini brings backgrounds from AWS, the Australian Cyber Security Centre, and the Department of Defence. Head of Engineering Brian Farnhill spent time at both AWS and Microsoft.

Drury, who built and exited email archiving platform Aftermail in the early 2000s before founding Xero, emphasised the depth of the communication problem. "It staggers me that organisational messaging remains this broken - with critical enterprise information fragmented across a myriad of consumer apps."

Corro is currently in development with a demonstration scheduled for 2026.

https://corro.co

Understanding Brotli PDF Compression

By Patrick Gallot

PDF has long relied on Flate for general-purpose data compression. Today, Brotli has emerged as a promising new approach that could significantly improve PDF compression ratios.

In today's world of instant global document sharing, PDF file size directly impacts performance. Whether you need faster downloads, reduced storage costs, or improved user experience, smaller PDF files can deliver big benefits.

PDF has long relied on Flate compression for general-purpose data compression, but Brotli has emerged as a promising new approach that could significantly improve compression ratios.

Originally developed by Google for web content compression, Brotli has proven its effectiveness in reducing web page load times. Now, as outlined in the PDF Association's recent article on Brotli compression, this advanced compression algorithm is making its way into the PDF specification, promising to squeeze PDF files more tightly than ever before.

PDF's Current Compression Landscape

PDF currently supports multiple compression methods, each optimized for different types of data. Among the general-purpose options - including RunLengthEncoding and LZW encoding - Flate compression is the most reliable for consistent results across various content types.

Flate compression, also known as Deflate, combines two compression techniques: the Lempel-Ziv algorithm (developed in the 1970s) and Huffman coding (from the 1950s) to create an efficient, lossless compression method.

Flate works as a dictionary coder: it scans through binary data searching for repeating patterns, then replaces these patterns with shorter references to entries in a dynamically built dictionary. During compression, Flate builds this dictionary from scratch for each data stream, embedding the dictionary information within the compressed output.

When decompressing, the algorithm reverses this process, using the embedded dictionary to reconstruct the original data.

This approach works well because it adapts to the specific characteristics of each data stream, but it also means that every compressed stream must carry its own dictionary information.

Brotli: A Smarter Approach to Compression

Brotli takes a different approach that builds upon the same foundation as Flate, but it adds something new: a predefined dictionary.

Rather than building a dictionary from scratch for each stream, Brotli starts with a dictionary that was synthesized using context modeling techniques to handle the most common data patterns found in web content (based on analysis conducted between 2013-2016).

This predefined dictionary is Brotli's secret weapon. Since it already contains patterns commonly found in text, HTML, CSS, JavaScript, and other web content, Brotli can immediately recognize and compress these patterns without needing to "learn" them first. For patterns not covered by the main dictionary, Brotli supplements this with a sliding window dictionary that detects other repeating patterns, similar to the way Flate operates.

The key advantage is efficiency: the more a data stream can leverage the predefined dictionary, the less dictionary information needs to be stored in the output, resulting in better compression ratios.

What This Means for PDF Compression

Brotli's effectiveness will largely depend on the type of content being compressed and how well it aligns with the patterns in its predefined dictionary. Content streams and font data are likely to be the biggest winners.

Text content, font definitions, and structured data within PDFs often contain patterns similar to web content—repeated keywords, common phrases, standardized formatting codes, and similar structural elements.

Image data is more complex. Losslessly compressed images contain data that's unlikely to match patterns in Brotli's web-optimized dictionary.

For these types of content, Brotli may perform similarly to Flate, or potentially less efficiently because it is attempting to match its predefined dictionary.

Strategic Implementation Approaches

Given these characteristics, two implementation strategies are possible:

■ Use Brotli as a drop-in replacement for Flate across all general-purpose compression needs. This approach banks on Brotli's average performance being better than Flate's. The simplicity of this approach makes it ideal for many use cases.

■ Implement logic to choose the optimal compression method for each data stream based on content type. This approach requires more sophisticated implementation but could yield maximum compression benefits by using Brotli where it excels and falling back to Flate or other methods where they perform better.

Looking to the Future

Brotli compression represents an exciting evolution in PDF optimization technology. By leveraging a predefined dictionary optimized for common content patterns, it offers the potential for significantly improved compression ratios, especially for text-heavy and structured content within PDF files.

Of course, in order to get the maximum benefit from using Brotli in PDFs, it has to be verified as an interoperable extension of PDF, formally standardized, and widely adopted.

Patrick Gallot, has been working with software developers and PDFs since 2000. At Datalogics, he is the lead technical support Engineer for the Adobe PDF Library and Datalogics PDF Java Toolkit, helping our customers resolve their complex questions and challenges. He is also active in the PDF community, presents at industry trade shows and events, and is a regular contributor to the Datalogics blog. Originally published here

Two Decades. One Vision.

Smarter Information Management

ELO Digital Office

AT THE ♥ OF YOUR BUSINESS

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

20 YEARS

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

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

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

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

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

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

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

Australian organisations address Shadow AI, but new risks are already emerging

By Tony Burnside, SVP and Head of APAC, Netskope

Shadow IT has always been a headache for security teams. Successive shifts in enterprise tech have always created blind spots and hidden behaviours that need to be discovered and secured. With the advent of SaaS apps, employees suddenly had a myriad of tools available to better collaborate and work. They quickly adopted them, often unaware of security approval processes, creating SaaS sprawl, shadow cloud, and significant changes in digital estates that some organisations are still grappling with.

The emergence of generative AI has brought similar challenges, with Shadow AI permeating most organisations. In 2024, a staggering 80% of Australian workers were using personal generative AI accounts at work. This is according to a report published by the Netskope Threat Labs team, which has been documenting the exposure of sensitive data through genAI usage in the workplace since 2023.

Their ongoing research, based on actual usage patterns from global and Australian businesses, shows that leaks of intellectual property, regulated data, or source code in genAI prompts are consistent across all geographies and industries, illustrating the data security risks of unmonitored and unsecured genAI usage.

But in a local cyber agenda dominated by pessimistic reports and cyber incidents, their most recent analysis delivers a glimmer of hope. In less than a year, the proportion of Australians using personal genAI accounts at work dropped sharply, from 80% to 55%.

This could be interpreted as a reduction in genAI use in the workplace, but this drop is a direct result of efforts by Australian organisations to centralise, gain visibility on, and secure genAI by deploying company-approved applications. It's a heartening finding that shows employees will adopt more secure behaviours when offered safe and easy-to-use alternatives.

Another sign of the growing maturity of AI security is Australian organisations' increased adoption of data loss prevention (DLP) tools, which rose to 41% from 32% last year.

DLP can inspect prompts and data in real time, and automatically block the transfer of sensitive information to unauthorised locations or contacts, acting as a safety net to human error.

Realtime user coaching tools can also support employees, presenting them with pop-ups when they attempt a risky action, and suggesting alternative, more secure options, or asking them to pause and justify or reconsider their action.

Research shows that a large majority (73% globally) choose not to proceed with their risky behaviour when presented with these coaching prompts.

LLMs, however, are only the tip of the AI iceberg, and AI-related cyber risk is becoming more multi-faceted as enterprise AI usage spreads and deployment models diversify. According to the Netskope Threat Labs researchers, 23% of Australian organisations are using on-premises LLM interfaces, and globally, 5.5% of organisations already have users running AI agents created using popular AI agent frameworks on-premises.

These platforms, which allow users to design and deploy AI models and agents within their organisation, are only going to grow in popularity in the upcoming months.

The appeal is the retention of data ownership, but these on-premises models often have very few inherent security features, and require proper configuration by security teams before they can be considered safe. Custom AI deployments, in most cases, also make use of open source resources, which can create AI supply chain security issues, some of which have already been reported.

Finally, most AI models or agents require direct access to enterprise data sources to train or complete their tasks, and without restricting their access levels and monitoring their operations, these systems could be over-permissioned and easily expose sensitive data.

Given the risks, security teams should make discovering AI deployments and eliminating any shadow AI within their organisation a priority. To achieve this, proactively communicating security protocols and principles for responsible AI use and development is an essential initial step.

From a security tools perspective, while DLP and real-time user coaching can help to an extent, they won't cover the multi-faceted risk surface created with AI projects, which can only be addressed with a comprehensive framework of security solutions.

With many security teams already managing large and fragmented ecosystems of security tools, the addition of yet more point security solutions should be avoided. Poorly integrated point solutions bring functionality overlap, gaps and duplicated management and operations burden.

Security platforms integrate multiple security tools into the same fabric and ensure they truly communicate and work in unison, delivering a more comprehensive and unified security approach.

Relevant platforms also enable security teams to manage and configure all their security tools from a single dashboard, and provide a unified and complete view of their data, network, traffic, and users across web, cloud, AI and private data centre environments.

Once AI initiatives are discovered, security teams can focus on ensuring all aspects of the project are secured, from the supply chain, to the stakeholders and data involved.

AI projects are only going to grow in complexity, and so will the challenge of securing the huge amounts of data they are going to interact with.

Australian organisations have been making great progress in tackling genAI risk, but need to keep their sights on the horizon, as new AI risks will continue to come quickly, and will require agility and speed in applying the appropriate security guardrails to prevent major incidents.

Tony Burnside is SVP and Head of APAC, Netskope

The Overlooked Risk in Retired IT Assets

Why End-of-Life Security Deserves a Front Seat in Data Protection Strategies

By Stuart Dahlenburg

In today's data-driven economy, organisations invest heavily in cybersecurity, compliance and digital transformation. Yet, one critical vulnerability continues to fly under the radar: the security risks associated with retired IT assets.

While leaders rightly focus on cybersecurity, as a frontline data protection measure, the conversation about IT Asset Disposition (ITAD) is often an afterthought. It shouldn't be. True data protection and operational resilience require ITAD to be an integral part of your strategy. After all, your reputation and compliance depend on it.

Recent research from Iron Mountain and Foundry* reveals a persistent disconnect. While many leaders acknowledge the security risks of retired IT assets, few are allocating the necessary budgets to address them. The research shows that 56% of IT leaders recognise the exposure risk from end-of-life hardware.

Yet, fewer than half ranked it among their top three data protection priorities, and most allocated less than 5% of security budgets to it.

3 drivers making ITAD a 2025 priority

There are three interconnected factors shifting ITAD to the forefront of business strategy in 2025: privacy, risk and sustainability. When we look at privacy, the regulatory landscape is growing tighter and accountability is moving up the chain.

For instance, Australian Privacy Principle 11.2 places a clear obligation on organisations to ensure that personal information is destroyed or de-identified when it is no longer required - and that includes the data from old laptops and network equipment.

In regulated sectors, rigorous lifecycle controls for decommissioning and disposal are no longer optional - they're an expected standard.

Beyond legal obligations, the financial and reputational risks are simply too high to ignore. A breach linked to improper disposal can be extremely costly, with average impacts estimated in the tens of millions of dollars.

The fallout isn't just financial; it erodes trust with customers and brings intense scrutiny from boards and regulators. It's easy to overlook ITAD because it's not the most visible part of the technology stack, but it is one of the most consequential.

Finally, sustainability is an undeniable part of business conversation and corporate reporting. As environmental, social, and governance (ESG) reporting requirements are phased in, a strong ITAD program is becoming essential for compliance.

It directly helps us track and reduce emissions, especially the waste generated in operations that falls under Scope 3. It also supports key circular economy goals, promoting the reuse and remarketing of assets and the recovery of materials.

As these disclosures mature, fluency in ITAD becomes a critical capability for any business leader.

What steps can Australian businesses take to lift ITAD readiness?

In a complex landscape of new regulations, evolving risk, and growing ESG requirements, secure ITAD is no longer a footnote - it's a critical control for safeguarding data and reputation.

While many leaders believe they understand the risks, the data paints a different picture. Iron Mountain's reveals a significant 79% of executives believe they understand the risks, yet nearly half admit their actions are inconsistent.

The remedy is to treat ITAD as a core enterprise risk control. A best-practice model for ITAD protects data at asset retirement, consistently enforces policy, and satisfies privacy and compliance requirements.

This approach also produces the necessary records for auditors and sustainability teams, ensuring a secure and compliant process from start to finish.

A practical guide for ITAD readiness

To help leaders get started, here's a simple checklist to assess your organisation's ITAD maturity. It covers the three essential dimensions you should measure yourself against:

Governance, compliance and auditability

- Is your ITAD policy documented and leadership-endorsed, with clear ties to regulation including, Australian Privacy Principle 11.2, CPS 234, and ISO 27001?
- Are roles, budgets and accountability clear, and are outcomes and value recovery consistently measured?
- Can you produce certificates and custody records for each asset within a week?

Security and chain of custody

- Can your providers demonstrate a secure chain of custody and data sanitisation on every job?
- Is your data sanitisation process aligned to NIST SP 800-88, and is it verified to match data sensitivity?

Sustainability and value

- Do you prioritise reuse and remarket before shredding, based on your risk profile?
- Does your ITAD program support ESG reporting, including Scope 3 Category 5 and other circular metrics?
- Are avoided emissions and waste reduction reported clearly and consistently?

If you cannot tick most of these boxes there's a gap that exposes your organisation to avoidable risk and regulatory scrutiny.

ITAD in Australia cannot be an afterthought; it must be a core part of your data protection strategy. If your program has gaps, you are exposed to unnecessary risk and scrutiny. Together let's make sure end-of-life device security is a priority, before it becomes the next front page headline.

Stuart Dahlenburg is General Manager, Asset Lifecycle Management, Iron Mountain.

** In Mar-Apr 2025 Foundry and Iron Mountain surveyed 317 IT-decision maker level respondents (Director and above titles in IT/ Networking/ Security, AI/ML, and Data/ Business Intelligence) from US,Australia, France and the UK working in the IT, Financial Services, Healthcare or Public Sector organisations.*

Automated DEWR Decisions 'Contrary to Law': Ombudsman

A Commonwealth Ombudsman investigation has found that automated government systems unlawfully cancelled income support payments for 9,642 job seekers over more than two years due to failures in updating computer systems after legislative changes.

The investigation into Australia's Targeted Compliance Framework (TCF) revealed that between April 2022 and July 2024, automated decisions cancelled welfare payments without the human discretion required under 2022 legislative amendments.

The Department of Employment and Workplace Relations (DEWR) and Services Australia failed to update their automated decision-making systems when Parliament changed the law to require consideration of individual circumstances before cancelling payments.

Commonwealth Ombudsman Iain Anderson found the agencies' actions were "contrary to law" and highlighted systemic failures in governance of automated decision-making processes.

"Automatic cancellation of vital income support for these job seekers is likely to have a profound, if not catastrophic impact," Anderson said in the report released this week.

The investigation uncovered multiple breakdowns in oversight and quality assurance. An independent review by Deloitte found the TCF computer system had become "increasingly unstable" with "volatility directly impacting compliance function operation."

DEWR's Secretary paused cancellation decisions in July 2024 after becoming aware of the unlawful automated processes. The department is now reviewing each affected case and making compensation payments.

The report made seven recommendations including establishing better consultation processes for legislative changes affecting automated systems, annual training for staff on administrative law requirements, and proactive identification of automation errors.

Of particular concern to compliance and risk managers, the investigation found that quality assurance activities failed to identify the unlawful cancellations for over two years. The automated system continued making decisions without human discretion despite legislative requirements.

The case highlights critical risks in automated decision-making systems, particularly the need to align technology with legal frameworks. The Ombudsman noted this was especially concerning given lessons from the Robodebt Royal Commission about automated processes affecting vulnerable people.

DEWR Secretary Natalie James accepted all recommendations and established a TCF Integrity Assurance Program with formal governance structures. The department committed to not resuming automated cancellations until systems comply with legal requirements.

The investigation also found DEWR failed to establish a Digital Protections Framework required by Parliament in 2022, more than three years after the legislative mandate.

Services Australia accepted recommendations for improved consultation processes, staff training, and ongoing assurance systems for automated decision-making compliance. The case serves as a warning for organisations implementing automated decision-making systems about the critical importance of governance, oversight, and alignment with regulatory requirements.

The full report is available here.

Cloud the Answer for Govt Legacy Woes

The Australian Government could unlock $A1.4 billion in annual productivity gains and cost savings by accelerating its migration to public cloud infrastructure, according to new research from consulting firm Mandala Partners, commissioned by Microsoft.

Their report reveals 71% of Commonwealth agencies still rely on legacy IT systems, and found that only 10% of total government IT procurement expenditure currently goes toward public cloud services, with smaller agencies leading adoption due to greater flexibility and less complex IT environments.

Accelerating cloud adoption by five years could deliver $A13.5 billion in total savings by 2035, representing a 13% reduction in IT expenditure compared to business-as-usual scenarios. The analysis shows government agencies could reduce their total IT budgets by between 7% and 28% over the decade through 2035, with larger agencies delivering the majority of savings.

The research identifies significant cybersecurity vulnerabilities in current infrastructure, with the Australian Government recording 163 data breaches in 2024 – the second highest of any sector. One legacy IT product still in use across government agencies has accumulated 1,653 known Common Vulnerabilities and Exposures in just five years since reaching end-of-life, with almost 30% categorised as severe.

Cloud migration could prevent an estimated 70 cyber incidents over the next decade, potentially saving $A178 million in breach-related costs through advanced security capabilities including automated threat detection, realtime monitoring, and AI-enhanced protection systems.

The report highlights substantial productivity benefits from cloud-enabled AI adoption, with an additional $A5 billion in productivity gains possible through widespread deployment of AI tools. The Australian Government's 2024 trial of Microsoft Copilot demonstrated these benefits, with participants saving up to one hour daily and 40% of saved time reallocated to higher-value activities including strategic planning and service improvement.

Cloud infrastructure would also support a 14% reduction in the government's carbon footprint and prevent 2.9 million hours of IT downtime through improved system reliability and automated failover mechanisms. The research identifies structural, cultural, and capability barriers preventing faster cloud adoption. Key challenges include budget frameworks that favour traditional short-term capital expenditure over operational expenditure models, risk-averse cultures that prioritise the status quo, and skills gaps with only 25% of agencies offering transformational learning programs for digital capabilities.

View the report here.

Preserving Legacy Through Pixels

Sir Henry Royce Foundation Embarks on Ambitious Digitisation Program

Trustee Howard Wraight's eye lit up like a kid with a new toy when the SHRF's new system was installed. "I am really impressed with how easy it all is to operate and how quickly and efficiently the scanner works."

In a quiet corner of Rowville, Victoria, a powerful transformation is underway - one that blends heritage with high-tech precision. The Sir Henry Royce Foundation Australia (SHRF), custodian of a rich archive chronicling the legacy of Henry Royce and the Rolls-Royce movement in Australia, has launched a comprehensive digitisation initiative designed to preserve its irreplaceable records for generations to come.

At the heart of this initiative is a complete scanning solution supplied and installed by Melbourne-based imaging specialists DocuVAN, known for their expertise in archival-grade digitisation across Asia Pacific.

The turnkey setup includes a high-resolution Fujitsu scanner, a robust HP workstation paired with dual 27" monitors for meticulous image review, two 14TB portable drives for redundant storage, and a dedicated uninterruptible power supply (UPS) to safeguard against data loss during power fluctuations.

"Records, photos and documents will be scanned at a quality often better than the original," remarked SHRF Chairman Brian Crump, underscoring the importance of digital fidelity in preserving fragile documents, photographs, and service records - many of which exist nowhere else.

"This is about confidence. Donors need to know their contributions will be professionally curated and protected."

Australia is one of three major archives for Sir Henry Royce and his legacy, which includes his whole philosophy around engineering excellence and innovation.

The UK, US and Australia, through charitable Foundations, hold the three major repositories of the considerable archival materials relating to Rolls-Royce and its history.

The Foundation's archival records collection ranges from books, magazines, brochures, RROCA and Branch historical documentation to service records from Australian Rolls-Royce and Bentley dealers. It includes sales and marketing catalogues dating back to the early 20th Century along with extensive technical information.

Trustee Howard Wraight said, "The variability and the diversity of material in our archive in terms of formatting, size and thickness is quite extensive posing a considerable challenge for high-speed scanning. The quantity of materials to be scanned is also a challenge, at a rough guess it's more than half a million pages that we have to scan.

"Do our records go back to the early 20th Century when the Rolls-Royce Silver Ghost was first made? Not extensively but the market in Victoria, Australia, and particularly Victoria's Western District, for Rolls-Royce in the early days pre-First World War was quite a big market. We were the third largest market for the company in the world. Obviously, other places have taken over since then, between the war era and certainly post-Second World War.

> The Foundation's archives are a treasure trove of automotive and aero-engine history, including rare publications, technical drawings, and other valuable records.

"We're not now a very big market for Rolls-Royce at all, but back then they were a very popular car, particularly amongst the wealthy landed gentry in the Western District of Victoria and also in the farming areas of NSW and Queensland, because they were just so well built and virtually indestructible compared to what else was available at the time. And some people were highly cashed up, so a number of cars came out during that period, of which there have been various publications documenting the early history of Rolls-Royce in Australia.

"Most of our extensive service records date from the late 50s through the 60s, 70s and into the 80s. Some of them, especially those coming from the 50s, are starting to look pretty sad and sorry, so it is really urgent that we get those preserved digitally so that they're not lost.

"We've also got more than just service records. We've got a very large number of photographs - more than 30,000 photos, which with this new high-resolution scanner we just put the photographs on the scanner and it automatically adjusts to the size and the shape and the material quality and just scans perfectly, so we can digitise all of our photographs as well."

The Foundation's archives are a treasure trove of automotive and aero-engine history, including rare publications, technical drawings, and other valuable records. Until now, these materials were vulnerable to environmental degradation and limited in accessibility. With DocuVAN's solution in place, SHRF can now digitise at scale, ensuring both preservation and discoverability.

"Everything's going to be scanned to a barcode," said Wraight. "So, the whole digitised record will sit behind a single barcode within the software that the scanner actually operates on, and then we're going to copy that across to an Excel database into which we put all the metadata. From searching the metadata, we can identify the original material via the barcode.

"As we load the record into the database, we identify the metadata that we then put in, which for a service record, for example, would be the manufacturer: Rolls-Royce or Bentley, the year of manufacturing, chassis number, the body type, engine number. We've got a whole list of all the things that we can enter in, plus also the origin of the service record as well - which company was doing the servicing. If we've got delivery dockets we can put all that in as well.

"So, if someone comes along and says, "Look, I've just bought a classic Rolls-Royce. I know the chassis number, or I know it's manufacturer and the body type and year of manufacture. Can you search for it?" We can do that.

"Once we get the process working really well, my vision is to select a couple of members of the Rolls-Royce Owners Club and ask them to volunteer to come and be trained, and we can just have people sitting there quietly scanning away the records on volunteer days. It's not going to take an inconceivably long period of time. I imagine we're going to get this all done, knocked over in a year or two."

"The solution provided by DocuVAN is exactly what the Foundation needed with our limited resources. It is cost effective, user friendly and with the high throughput scanning capabilities it has put the feasibility of digitally preserving our vast collection of historical materials well within our reach, something that did not seem possible until now."

DocuVAN's tailored approach reflects a growing trend in heritage institutions embracing digital transformation. Their solutions cater to a wide range of archival needs - from microfilm and fragile records to oversized artwork and books.

The digitisation program also aligns with SHRF's broader strategic vision, which includes publishing an annual report in 2026 and expanding public access to its collections via electronic platforms.

SHRF Trustee Clare Gordon said, "I am confident that this approach is sustainable and allows the Foundation to be in control of its asset collection material appropriately, capture the documentation in a safe and secure manner and prepare the document collection to be able to be presented for electronic access."

https://www.henryroycefoundation.com

For more on DocuVAN's scanning solutions, visit DocuVAN's official site.

Original 1914 Rolls-Royce Silver Ghost Catalogue

Firms Accelerate IDP Projects Amid Skills Gap: Survey

Nearly two-thirds of large enterprises are accelerating intelligent document processing projects as artificial intelligence transforms how organisations handle unstructured documents, according to new research.

The survey of 600 US and European organisations found 65 per cent are actively considering or implementing new IDP initiatives, with productivity gains rather than workforce reduction driving adoption.

However, the research revealed a paradox in digital transformation efforts: 61 per cent of IDP processes still involve paper, and 48 per cent of respondents expect their paper use to increase.

"Despite significant digital transformation efforts, paper remains entrenched in business processes," the study found. European firms anticipate a 50 per cent drop in paper use within the next year, while US companies expect only a 30 per cent reduction.

The research, conducted by the Association for Intelligent Information Management with Deep Analysis and co-sponsored by SER Group, surveyed organisations with revenues exceeding $US10 million and more than 500 employees.

Front-office applications lead growth

The study identified a shift from traditional back-office functions like invoice processing to front-office use cases including HR files, contracts, licences and permits, and customer onboarding.

Respondents cited reduced processing time as the biggest benefit (50 per cent), rather than headcount reductions (30 per cent), positioning IDP as a productivity enhancer rather than a workforce reduction tool.

However, implementation faces significant barriers. More than 50 per cent of respondents reported shortages of technological expertise and process redesign skills, while data security and privacy concerns topped the list of adoption challenges.

Skills shortage hampers adoption

The research highlighted the need for organisations to invest in change management and cross-functional training to maximise return on investment from IDP implementations.

Integration challenges ranked as the second-biggest barrier, underscoring requirements for interoperable standards and embedded IDP capabilities within existing enterprise systems.

The study also found that generative AI-powered research has overtaken traditional references and analyst reports as the leading method for selecting IDP solutions, though proof-of-concept testing and industry-specific expertise remain essential in purchasing decisions.

The research indicates that 78 per cent of surveyed organisations already use some form of AI, confirming IDP as an established enterprise use case for artificial intelligence adoption.

WA marks progress on Electronic Medical Record

Western Australia has completed the first phase of its rollout of an Electronic Medical Record (EMR) Program, marking a major milestone in the state's healthcare modernisation efforts. The WA Government announced that its Electronic Medical Record (EMR) Program has successfully implemented Single Sign On (SSO) and Digital Medical Record (DMR) systems across all WA Health facilities, part of a $A247 million investment.

The completion in August 2025 represents the culmination of a multi-year digital transformation project that began with the SSO rollout in February 2023. The systems are now operational across Metropolitan Perth and regional areas including the Mid West, Kimberley, Wheatbelt, Great Southern, South West, and Pilbara.

The SSO system, is used by 27,000 clinicians across WA Health with 348,000 logins per week, reducing time clinicians spend logging into their systems, allowing them to focus more on patient care. The DMR is digitising paper-based patient records. As of July 2025, the system has digitised 44 million documents and serves approximately 13,000 WA Health staff on a typical weekday.

The SSO system has simplified access for clinicians, allowing them to access around 90 clinical applications with a single card tap. This reduces the time healthcare workers spend on administrative tasks and allows greater focus on patient care.

Meanwhile, the DMR system has replaced paper-based patient records, providing clinicians with realtime access to complete, up-to-date patient information across all WA Health facilities statewide.

Health Minister Meredith Hammat emphasized the significance of the achievement for the state's vast healthcare network.

"This significant milestone is about modernising and improving healthcare right across our State," Hammat said.

"Western Australia is geographically the largest health jurisdiction in the world, and these technologies help to bring us closer together and improve patient care."

The $247 million investment includes $104 million allocated in the 2024-25 State Budget, with the government now moving toward planning and procurement of a comprehensive Electronic Medical Record system as the program's next phase. "These innovations will help ensure that WA Health remains at the forefront of modern healthcare delivery for our community."

New Study Raises Serious Privacy Concerns over AI Assistants

Popular generative AI web browser assistants are collecting and sharing sensitive user data, such as medical records and social security numbers, without adequate safeguards, finds a new study led by researchers from University College London (UCL) and Mediterranea University of Reggio Calabria.

The study is claimed to be the first large-scale analysis of generative AI browser assistants and privacy. It uncovered widespread tracking, profiling, and personalisation practices that pose serious privacy concerns, with the authors calling for greater transparency and user control over data collection and sharing practices.

The researchers analysed nine of the most popular generative AI browser extensions, such as ChatGPT for Google, Merlin, and Copilot (not to be confused with the Microsoft app of the same name). These tools, which need to be downloaded and installed to use, are designed to enhance web browsing with AI-powered features like summarisation and search assistance, but were found to collect extensive personal data from users' web activity.

Analysis revealed that several assistants transmitted full webpage content – including any information visible on screen – to their servers. One assistant, Merlin, even captured form inputs such as online banking details or health data.

> "... these AI browser assistants operate with unprecedented access to users' online behaviour in areas of their online life that should remain private."

Extensions like Sider and TinaMind shared user questions and information that could identify them (such as their IP address) with platforms like Google Analytics, enabling potential cross-site tracking and ad targeting.

ChatGPT for Google, Copilot, Monica, and Sider demonstrated the ability to infer user attributes such as age, gender, income, and interests, and used this information to personalise responses, even across different browsing sessions.

Only one assistant, Perplexity, did not show any evidence of profiling or personalisation.

Dr Anna Maria Mandalari, senior author of the study from UCL Electronic & Electrical Engineering, said: "Though many people are aware that search engines and social media platforms collect information about them for targeted advertising, these AI browser assistants operate with unprecedented access to users' online behaviour in areas of their online life that should remain private.

"While they offer convenience, our findings show they often do so at the cost of user privacy, without transparency or consent and sometimes in breach of privacy legislation or the company's own terms of service.

"This data collection and sharing is not trivial. Besides the selling or sharing of data with third parties, in a world where massive data hacks are frequent, there's no way of knowing what's happening with your browsing data once it has been gathered."

For the study, the researchers simulated real-world browsing scenarios by creating the persona of a 'rich, millennial male from California', which they used to interact with the browser assistants while completing common online tasks.

This included activities in both the public (logged out) space, such as reading online news, shopping on Amazon or watching YouTube videos. It also included activities in the private (logged in) space, such as accessing a university health portal, logging into a dating service or accessing pornography. The researchers assumed that users would not want this activity to be tracked due to the data being personal and sensitive.

During the simulation the researchers intercepted and decrypted traffic between browser assistants, their servers and third-party trackers, allowing them to analyse what data was flowing in and out in real time. They also tested whether assistants could infer and remember user characteristics based on browsing behaviour, by asking them to summarise the webpages then asking the assistant questions, such as 'what was the purpose of the current medical visit?' after accessing an online health portal, to see if they had retained personal data.

The experiments revealed that some assistants, including Merlin and Sider, did not stop recording activity when the user switched to the private space as they are meant to.

The authors say the study highlights the urgent need for regulatory oversight of AI browser assistants in order to protect users' personal data. Some assistants were found to violate US data protection laws such as the Health Insurance Portability and Accountability Act (HIPAA) and the Family Educational Rights and Privacy Act (FERPA) by collecting protected health and educational information.

The study was conducted in the US and so compatibility with UK/EU data laws such as GDPR was not included, but the authors say this would likely be a violation in the EU and UK as well, given that privacy regulations in those places are more stringent.

The authors recommend that developers adopt privacy-by-design principles, such as local processing or explicit user consent for data collection.

Dr Aurelio Canino, an author of the study from UCL Electronic & Electrical Engineering and Mediterranea University of Reggio Calabria, said: "As generative AI becomes more embedded in our digital lives, we must ensure that privacy is not sacrificed for convenience. Our work lays the foundation for future regulation and transparency in this rapidly evolving space."

Government AI Survey Reveals Adoption Gaps

While the Australian Government last month released a set of AI Technical Standards, significant capability gaps exist across public sector agencies, with some yet to begin adoption despite mounting risks from technological inaction.

Digital Transformation Agency Deputy CEO Lucy Poole told the Tech in Gov 2025 conference that the AI standards establish critical frameworks for data interoperability, algorithmic transparency and model validation across government AI systems.

"In government, trust is everything. One misstep can stall progress and erode the social license we rely on. And once it's lost, it's a long road back," Poole said during Tuesday's keynote address.

The standards emerge as the DTA's AI Accountable Officials Survey of 80 agencies revealed adoption disparities. While most agencies are actively trialling AI tools and building staff capability, some have yet to begin implementation in what Poole described as a high-risk position.

"In the rapidly changing landscape of technology, standing still is not a neutral position; it exposes organisations and the communities they serve to significant risk," she warned.

Legacy Infrastructure Creates Major Barriers

Poole identified legacy systems as the primary impediment to AI adoption, citing outdated infrastructure, poor data management and compatibility issues that make integration challenging and costly.

"Legacy systems hinder AI adoption due to outdated infrastructure, poor data management, and compatibility issues, making integration challenging. Modernising these systems is costly and time-consuming," she said.

The resource drain affects capability development, with legacy system maintenance "diverting resources from AI initiatives, slowing down the adoption process."

Smaller agencies face particular challenges, reporting gaps in skills, resources and funding alongside concerns about data privacy and legal complexity. The survey findings highlight sector-wide capability issues that could undermine progress and public trust.

Standards Enable Platform-Scale Implementation

The AI Technical Standards represent a shift from experimental pilots toward enterprise-scale deployment, Poole explained. Drawing on a CSIRO Data61 metaphor, she emphasised that governance frameworks enable rather than constrain innovation.

"Brakes don't slow us down - they help us go faster," she said, quoting Dr Liming Zhu. "The safeguards we build - our standards, governance, and assurance tools - aren't barriers. They're enablers."

The standards were developed through cross-government collaboration and embed transparency, accountability and safety requirements. They aim to support agencies moving "from pilots to platforms - from isolated experiments to whole-of-government capability."

Pool showcased The National Library of Australia's oral history digitisation as exemplifying scalable AI deployment, with systems transcribing one hour of audio in 90 seconds. "That means decades of voices—stories, lived experiences—are now searchable and accessible to researchers, educators, and the public," Poole said.

Risk Management and Public Trust

The presentation acknowledged that weak implementation by individual agencies could compromise sector-wide progress and public confidence.

"One weak link in our chain can have serious consequences - compromising not only the effectiveness of our systems, but also the trust of the public," Poole warned.

The standards framework aims to maintain public trust while enabling innovation, with governance mechanisms providing "the guardrails, the map, and the brakes we need to navigate complexity with confidence."

Poole concluded by referencing Mo Gawdat's observation that AI adoption represents more than technological change: "This isn't just a story about technology. It's about us—human nature, ethics, and how we choose to handle this powerful tool."

The DTA serves as the Australian Government's adviser for whole-of-government digital and ICT strategies, policies and standards, including procurement oversight.

What is Data Custodianship?

By Nicola Askham

What is Data Custodianship? It's a question that seems straightforward, but sometimes in Data Governance, definitions can vary depending on who you ask. And this is okay because some terms might not suit the culture or structure of your organisation. It's always best to prioritise what will work for your organisation specifically.

However, it is something to be aware of, particularly when 'googling' terms online or you're new to an organisation. Do not assume that when colleagues say Data Custodianship, they mean what I'm about to tell you. I would say the best thing to do if somebody uses a term like Data Custodian is to ask them what they mean by that and who that is. This clears things up from the start. You'll be on the same page as everyone else and can have useful conversations.

When I talk about Data Custodians, I'm referring to the IT department. I often talk about the importance of business in taking ownership in managing and understanding data but, in this instance, IT remains a crucial player.

The role of Data Custodian differs from other roles in Data Governance, such as data Owners and Data Stewards etc. With these roles, I'd always recommend that you go and find named individuals. However, the opposite is true when it comes to Data Custodians because generally, I would say the whole of your IT department are Data Custodians. There are so many different areas of expertise and disciplines within an IT department that no one person will know absolutely everything about that system to be the Data Custodian for it. So I usually just say that IT are the Data Custodians for all the data that's held on IT supported systems at your organisation.

The responsibility of IT, as Data Custodians, lies in maintaining data on systems in line with the requirements of the business. This involves tasks such as data maintenance, migration, aggregation and transformation - all guided by business needs.

The misconception that IT own the data

FIt's important to realise that IT does NOT own the organisation's data. Yes, it is on their systems, and they have the expertise that the business side of the organisation may not have, but IT shouldn't be expected to work out what to do with the data.

Before the introduction of formal Data Governance, businesses often rely on IT to make business data-related decisions, but this can lead to them being blamed for things. I think is unfair because sometimes they're just doing the best they can with poor requirements from the business.

Data Custodianship helps to clarify roles and responsibilities within the IT department. When I work with IT departments, they are pleased with this Data Governance role as it gives them very clear business requirements and named business people to go to make decisions about data.

So this is a really good way of starting to break down some silos and get the business to understand what happens to the data when it's on systems. IT has always played the role of Data Custodian but, without Data Governance, they've perhaps done it without the input they needed from the business.

So having a Data Governance framework in place and identifying IT as Data Custodians is a really good way to start improving communications and making consistent, holistic decisions about data.

Understanding Data Custodianship is essential for establishing effective Data Governance practices. By recognising the roles of both the business and IT, organisations can foster collaboration, enhance data quality and make informed decisions that align with business objectives.

Originally published on www.nicolaaskham.com.

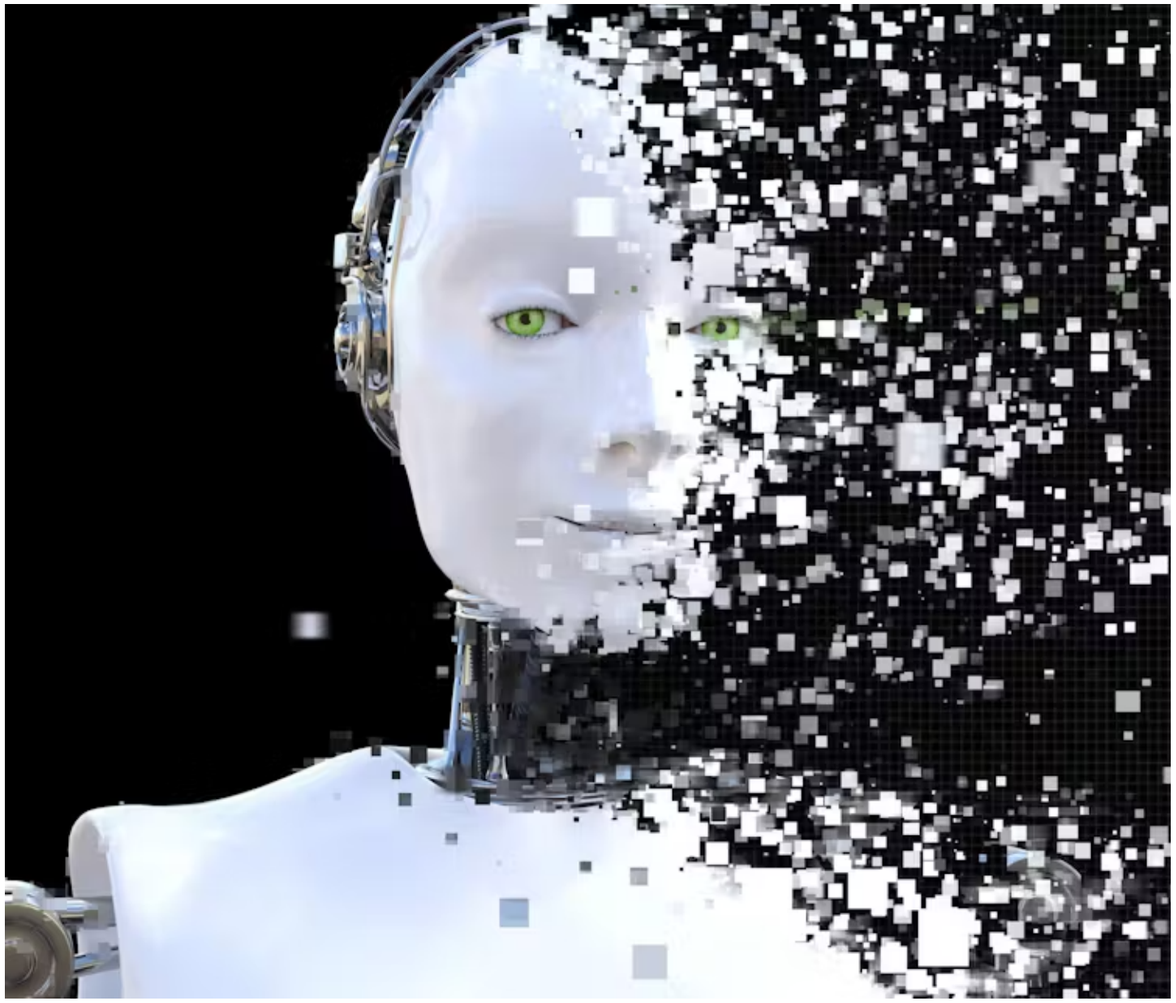

Why OpenAI's solution to AI hallucinations would kill ChatGPT

By Wei Xing, University of Sheffield

OpenAI's latest research paper diagnoses exactly why ChatGPT and other large language models can make things up – known in the world of artificial intelligence as "hallucination". It also reveals why the problem may be unfixable, at least as far as consumers are concerned. The paper provides the most rigorous mathematical explanation yet for why these models confidently state falsehoods. It demonstrates that these aren't just an unfortunate side effect of the way that AIs are currently trained, but are mathematically inevitable.

The issue can partly be explained by mistakes in the underlying data used to train the AIs. But using mathematical analysis of how AI systems learn, the researchers prove that even with perfect training data, the problem still exists.

The way language models respond to queries – by predicting one word at a time in a sentence, based on probabilities – naturally produces errors. The researchers in fact show that the total error rate for generating sentences is at least twice as high as the error rate the same AI would have on a simple yes/no question, because mistakes can accumulate over multiple predictions.

In other words, hallucination rates are fundamentally bounded by how well AI systems can distinguish valid from invalid responses. Since this classification problem is inherently difficult for many areas of knowledge, hallucinations become unavoidable.

It also turns out that the less a model sees a fact during training, the more likely it is to hallucinate when asked about it. With birthdays of notable figures, for instance, it was found that if 20% of such people's birthdays only appear once in training data, then base models should get at least 20% of birthday queries wrong.

Sure enough, when researchers asked state-of-the-art models for the birthday of Adam Kalai, one of the paper's authors, DeepSeek-V3 confidently provided three different incorrect dates across separate attempts: "03-07", "15-06", and "01-01". The correct date is in the autumn, so none of these were even close.

The evaluation trap

More troubling is the paper's analysis of why hallucinations persist despite post-training efforts (such as providing extensive human feedback to an AI's responses before it is released to the public). The authors examined ten major AI benchmarks, including those used by Google, OpenAI and also the top leaderboards that rank AI models. This revealed that nine benchmarks use binary grading systems that award zero points for AIs expressing uncertainty.

This creates what the authors term an "epidemic" of penalising honest responses. When an AI system says "I don't know", it receives the same score as giving completely wrong information. The optimal strategy under such evaluation becomes clear: always guess.

The researchers prove this mathematically. Whatever the chances of a particular answer being right, the expected score of guessing always exceeds the score of abstaining when an evaluation uses binary grading.

The solution that would break everything

OpenAI's proposed fix is to have the AI consider its own confidence in an answer before putting it out there, and for benchmarks to score them on that basis. The AI could then be prompted, for instance: "Answer only if you are more than 75% confident, since mistakes are penalised 3 points while correct answers receive 1 point."

The OpenAI researchers' mathematical framework shows that under appropriate confidence thresholds, AI systems would naturally express uncertainty rather than guess. So this would lead to fewer hallucinations. The problem is what it would do to user experience.

Consider the implications if ChatGPT started saying "I don't know" to even 30% of queries – a conservative estimate based on the paper's analysis of factual uncertainty in training data. Users accustomed to receiving confident answers to virtually any question would likely abandon such systems rapidly.

I've seen this kind of problem in another area of my life. I'm involved in an air-quality monitoring project in Salt Lake City, Utah. When the system flags uncertainties around measurements during adverse weather conditions or when equipment is being calibrated, there's less user engagement compared to displays showing confident readings – even when those confident readings prove inaccurate during validation.

The computational economics problem

It wouldn't be difficult to reduce hallucinations using the paper's insights. Established methods for quantifying uncertainty have existed for decades. These could be used to provide trustworthy estimates of uncertainty and guide an AI to make smarter choices.

But even if the problem of users disliking this uncertainty could be overcome, there's a bigger obstacle: computational economics. Uncertainty-aware language models require significantly more computation than today's approach, as they must evaluate multiple possible responses and estimate confidence levels. For a system processing millions of queries daily, this translates to dramatically higher operational costs.

More sophisticated approaches like active learning, where AI systems ask clarifying questions to reduce uncertainty, can improve accuracy but further multiply computational requirements. Such methods work well in specialised domains like chip design, where wrong answers cost millions of dollars and justify extensive computation. For consumer applications where users expect instant responses, the economics become prohibitive.

The calculus shifts dramatically for AI systems managing critical business operations or economic infrastructure. When AI agents handle supply chain logistics, financial trading or medical diagnostics, the cost of hallucinations far exceeds the expense of getting models to decide whether they're too uncertain. In these domains, the paper's proposed solutions become economically viable – even necessary. Uncertain AI agents will just have to cost more.

However, consumer applications still dominate AI development priorities. Users want systems that provide confident answers to any question. Evaluation benchmarks reward systems that guess rather than express uncertainty. Computational costs favour fast, overconfident responses over slow, uncertain ones.

Falling energy costs per token and advancing chip architectures may eventually make it more affordable to have AIs decide whether they're certain enough to answer a question. But the relatively high amount of computation required compared to today's guessing would remain, regardless of absolute hardware costs.

In short, the OpenAI paper inadvertently highlights an uncomfortable truth: the business incentives driving consumer AI development remain fundamentally misaligned with reducing hallucinations. Until these incentives change, hallucinations will persist.

Wei Xing is Assistant Professor, School of Mathematical and Physical Sciences, University of Sheffield, This article is republished from The Conversation under a Creative Commons license. Read the original article.

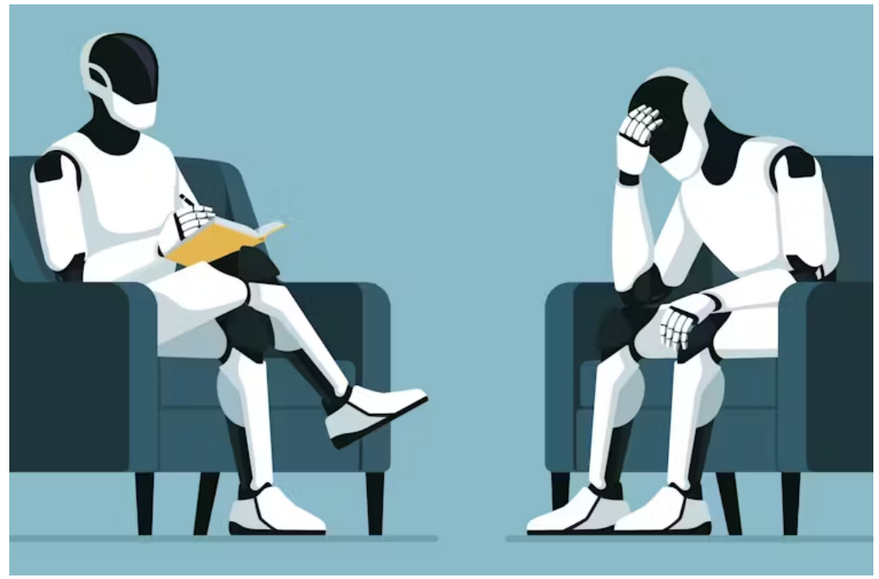

'Have as many crazy guesses as you like.' ElenaBs/Alamy

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Knowledge Agent Preps SharePoint for AI

Microsoft has introduced Knowledge Agent for SharePoint, an AI-powered content management assistant that automates tagging and classification to eliminate manual effort and inconsistency.

"Your AI suite is only as powerful as the content it's built on," said John Mighell, Microsoft's Director, Product Marketing, Microsoft SharePoint and OneDrive in a blog post announcing the launch. "If your content isn't ready, your AI tools won't be either."

Knowledge Agent aims to solve "the most pressing content management challenges – like content readiness for AI, discoverability and freshness, manual governance processes, and content creation bottlenecks."

Knowledge Agent intelligently tags and classifies files with auto-filled metadata, with Copilot and agents able to reason over this data starting next month. The system suggests metadata columns based on content analysis, eliminating manual tagging efforts that have traditionally burdened information management teams.

The AI-powered system creates automated views that sort, filter, and group documents based on metadata, helping users quickly find relevant content such as "policies expiring in 2026" or "contracts grouped by client."

Microsoft positions the technology to help SharePoint "drive real business impact, not just deliver answers," according to the announcement.

For organisations struggling with compliance requirements, the system provides automated site maintenance capabilities. Knowledge Agent analyses search behaviour to detect content gaps, automatically fixes broken links, and enables administrators to retire inactive pages in just a few clicks.

The platform enables business users to create automated workflows using natural language commands. Users can describe requirements such as "email me when invoices over $500 are added" and the agent builds the workflow without technical expertise.

"No more spending hours tagging manually," Mighell noted, highlighting how the system "suggests autofill columns based on content and user input, ensuring consistent organisation and better discoverability."

Administrators can control site participation through PowerShell commands, with options to enable Knowledge Agent across all sites, no sites, or all sites except selected ones. The exclusion list is limited to 100 sites maximum.

The system requires tenant-level opt-in initially, with individual site opt-in capabilities planned from November 1, 2025.

Knowledge Agent appears as a floating button in SharePoint's lower-right corner, providing context-aware suggestions based on user roles and current location within the platform. Site owners see options for site improvement, while document library editors access organisation tools and workflow creation features.

According to Mighell, Knowledge Agent "transforms web content creation into a dynamic, multi-turn experience" and "acts as your co-author for exceptional web experiences."

The interface adapts to user permissions, ensuring appropriate functionality access while maintaining security controls essential for regulated industries.

Microsoft expects general availability in early 2026, with additional licensing details to be communicated before that release.

Archive360 and Neev Data Tackle SAP Data Silos

Archive360 and Neev Data have announced a strategic partnership to combine their platforms to deliver a single, governed repository for structured ERP records, unstructured documents, and all communication channels.

The solution integrates Neev Data's SAP-certified connectors with Archive360's cloud governance platform. Neev's software extracts and normalises data from SAP and other legacy applications. The data is then managed within the Archive360 platform, which applies indexing, classification, and retention policies.

This allows compliance and legal teams to search and manage ERP records alongside emails and chat logs from a single console.

The partnership addresses a critical challenge facing organizations today: achieving a true «single pane» view of archived data. Most compliance teams must juggle separate tools for SAP and other ERP system content, as well as for email, chat, and other communications, all of which increases costs, complexity and risk.

According to the vendors, the unified platform can help decommission legacy systems when migrating to new environments like SAP S/4 HANA, while retaining access to historical data. The firms also claim that unifying this data makes it accessible for AI and analytics programs.

Neev's connectors extract SAP tables, documents, and attachments in an open format at scale and hand them off to Archive360's platform for indexing, classification, and retention management.

Archive360's engine applies a single set of compliance rules across both SAP and other ERP archived data along with communication data, ensuring consistent policy enforcement and providing holistic search capabilities.

https://www.archive360.com

https://www.neevdata.com

Affinda Claims Agentic IDP Breakthrough

Enterprise document processing has long suffered from a fundamental flaw: systems that can't learn from their mistakes without extensive retraining. Melbourne-based Affinda claims to have solved this persistent industry challenge with the launch of its new agentic AI intelligent document processing (IDP) platform featuring "persistent model memory."

The breakthrough addresses what Affinda positions as a critical gap in the current IDP market, where organizations must choose between accuracy, speed, and adaptability. While the company claims competitors require weeks or months of fine-tuning with dozens of examples, Affinda's platform is able to recall any corrections that were required for processing individual document types and automatically apply the corrections the next time a similar document is presented.

This represents a significant departure from the approaches used by established IDP providers, according to Andrew Bird, Affinda's Head of AI, who emphasizes that the platform's ability to remember and apply corrections without human intervention sets it apart from traditional LLM-powered document processing tools that rely on static prompts.

"If a human corrects a model in the Affinda Platform for a particular document, this correction is added to the platform's model memory and it instantly learns not to make that mistake again. Systems that rely on fine-tuning smaller models will likely require some number of examples (usually 20+) before the model is likely to learn from its mistakes in the next training run. It should come as no surprise that an approach that leverages frontier models is becoming table stakes in IDP," said Bird.

(A "frontier model" refers to the most advanced, state-of-the-art large language models that represent the cutting edge of AI capabilities, e.g. GPT-4o and o1 from OpenAI, Claude 4 (Opus and Sonnet) from Anthropic and Gemini Ultra/Pro from Google DeepMind.)

The technology addresses persistent challenges in digital transformation initiatives where document-heavy processes remain manual bottlenecks. Traditional approaches typically require extensive training periods and ongoing maintenance, while standalone large language models face performance limitations.

The Affinda platform operates from data centres in Sydney, Frankfurt and Oregon, offering both cloud-based and self-hosted deployment options. Affinda is ISO27001 certified and GDPR compliant, and will soon be SOC2 accredited, meeting enterprise security and privacy requirements.

Anthony England, Head of Growth, Affinda, said: "Until now, enterprises that wanted to automate their document-heavy processes were forced to compromise on accuracy, time to value or flexibility. Now they can get all three. Affinda's platform can accurately extract any information from any document, in any format, fast. It's world class, and our customers are blown away by how quickly they can get set up and how powerful the platform is.

"Using the Affinda platform, organizations can now build and customize document processing models themselves in just a few clicks, paying only for what they use. By removing the need for extensive DIY and in-house R&D, we're taking the guesswork out of automating document workflows, and we're giving time back to organizations to focus on core projects and innovation."

The Affinda platform combines frontier large language models hosted on AWS and Azure with internally fine-tuned smaller models for pre- and post-processing. The system uses Microsoft Azure's OCR engine and Elasticsearch for vector searches in its retrieval-augmented generation pipeline.

Affinda claims organizations can build and deploy models in minutes rather than months, achieving what the vendor states is over 99% accuracy across all document types. The platform supports 56 languages and integrates with more than 400 enterprise systems.

Two enterprise customers have detailed their implementations. Insurance provider StateCover Mutual processes over 300,000 documents annually across 80 document types using the platform. Global logistics company Northline reports processing 120,000 proof-of-delivery and related documents with 82% straight-through processing rates.

"Automating validation and document handling across all 13 depots has reduced manual effort, errors and delays," said Jorg Both, Northline's head of business systems.

Regarding competitive positioning, Affinda claims its approach outperforms established players including Microsoft's Azure AI Document Intelligence and Google's recently released Document AI platform.

Microsoft's Azure service relies on fine-tuning smaller models rather than LLMs, according to Affinda, while Google's new platform uses LLM fine-tuning with what the vendor characterizes as slower learning cycles compared to its instant model memory updates.

The platform targets organizations managing compliance-heavy document workflows, offering consumption-based pricing and a free trial with access to all features.

https://www.affinda.com

Artificio Integrates PDF Creation into Workflow

California-based Artificio has expanded its AI-powered document processing platform to include PDF generation capabilities, enabling organisations to create new documents from data extracted through existing processing workflows.

The addition allows enterprises to process incoming documents using AI agents for data extraction and classification, then automatically generate PDF reports, contracts or presentations based on extracted insights within the same platform.

The new feature includes drag-and-drop editing tools, automated data integration from processed documents, and template automation for generating compliance documents and reports based on extracted information.

Artificio's platform combines document intake and classification, data processing and analysis, workflow automation, and now document creation within a single system. The company positions this as addressing enterprise demand for unified document intelligence solutions rather than separate processing and creation tools.

Applications span financial services for loan processing and compliance reporting, healthcare for patient record analysis and billing documentation, legal firms for contract processing and client communications, and supply chain organisations for purchase order management and vendor communications.

The cloud-based platform maintains API connectivity for integration with existing enterprise systems and includes digital signature capabilities.

https://artificio.ai

New Tools Identify Shadow AI Risks

Data security platform BigID has launched four new AI governance capabilities designed to help organisations discover unauthorised AI models and control sensitive data use in artificial intelligence systems. The New York-based company announced the suite is targeting what it claims are growing risks from "shadow AI" – unauthorised AI models operating without security oversight.

The capabilities include Shadow AI Discovery to uncover rogue models, Data Labeling for AI to classify appropriate datasets, Data Cleansing for AI to remove sensitive information, and what BigID calls the industry's first Prompt-Based Classification system using natural language.

BigID's Shadow AI Discovery automatically identifies unmanaged AI models across cloud and collaboration platforms, while flagging personal or regulated training data. The system aims to provide visibility into AI usage that traditional security tools miss, integrating across model repositories, developer tools, cloud platforms and collaboration environments.

The capability goes beyond discovery to enable direct enforcement actions. Security teams can trigger policy enforcement, restrict risky access and launch remediation workflows within the BigID platform. The system correlates models to underlying datasets and maps out user activity patterns to show who is using what AI tools, where and how across the enterprise environment.

The Data Labeling for AI feature helps organisations classify and tag data for appropriate AI use through usage-based labels. BigID provides out-of-the-box classifications including "AI-approved," "restricted," and "prohibited," while allowing organisations to create custom labels aligned with internal risk frameworks and regulatory requirements.

The system supports both structured and unstructured data across cloud, software-as-a-service and collaboration environments. It aims to enforce usage policies early in data pipelines before information reaches AI models, combining classification with policy enforcement and remediation workflows to prevent sensitive or high-risk data from entering large language models, copilots and retrieval-augmented generation systems.

Data Cleansing for AI focuses on removing or tokenising sensitive information at scale before it enters generative AI tools and large language models. The capability works across both structured and unstructured data formats, including emails, PDFs, collaboration files and databases, to prevent confidential data from being embedded into model outputs or leaked in prompts.

The system forms part of BigID's broader Secure Data Pipeline solution, working alongside other capabilities including GenAI Catalog, Search and Safe-for-AI Labeling features. Security teams can apply policy-based controls and continuously reduce exposure across AI initiatives, with the goal of strengthening generative AI pipeline security through pre-cleansed, policy-compliant datasets.

BigID's Prompt-Based Classification system replaces traditional rule-based data classification with what the company claims is an industry-first natural language interface. Users can describe sensitive data in plain English, paste regulatory language, or articulate AI policy requirements, with BigID's AI-powered engine automatically generating classification logic.

The system aims to address limitations of traditional classification tools that rely on technical rules, pattern libraries and complex configurations often inaccessible to non-technical teams. The capability scales discovery across cloud, software-as-a-service, data lakes and file systems without manual rule creation, while promising reduced false positives through context-aware intelligence that understands use case and risk factors.

https://www.bigid.com

AI Engine Automates Access Control

Delinea has launched Iris AI, an artificial intelligence engine that automates identity security decisions and auditing processes for enterprise IT systems.

The cloud-native platform evaluates user behaviour, device location, and policy alignment to grant or deny access requests in realtime, while automatically analysing recorded sessions to detect suspicious activities and policy violations.

Identity security has become increasingly complex as organisations adopt hybrid cloud environments and expand their digital footprints, creating challenges for IT teams managing thousands of user accounts and machine identities across multiple platforms.

"In today's fast-paced, hybrid cloud environments, teams are constantly spinning up new apps, tools and services, often without security oversight," said Jackie McGuire, security practice lead and principal analyst at theCUBE Research. "With Delinea Iris AI, there's real value in transforming shadow IT into a security onboarding pipeline, giving teams visibility and control across every identity, no matter where it lives or how it's managed."

The platform includes two primary capabilities:

- Authorization powered by Delinea Iris AI builds context in real-time by evaluating user behaviour, business justification, device, location, and policy alignment to intelligently triage risk for every access request without slowing productivity. Dynamically adjust access as user context changes, offering clear, evidence-based reasoning and a complete audit trail.

- Auditing powered by Delinea Iris AI analyzes recorded sessions, detecting elevated privileges, failed authorizations, deletions, file transfers, unusual Secrets usage, and more. It highlights elevated commands and risky behaviours in seconds, delivering an evidence-based summary and heatmap of suspicious activity within each session. This provides instant, actionable insights to quickly identify threats, investigate issues, and stop bad actors before the damage is done.

Delinea operates globally with a focus on centralised authorisation systems.

https://delinea.com

Cloudflare Platform blocks Shadow IT

Cloudflare has integrated its Cloud Access Security Broker (CASB) platform with three major generative AI tools to address growing enterprise security concerns over workplace AI adoption.

The connectivity cloud company claims to be the first CASB provider to integrate directly with ChatGPT Enterprise, Claude by Anthropic, and Google Gemini, enabling realtime monitoring and control of employee AI usage.

The integrations allow Cloudflare's CASB to continuously scan these platforms for potential data security risks, providing automated alerts when sensitive corporate information may be exposed through AI interactions.

Notably absent from Cloudflare's integrations is Microsoft Copilot, despite its widespread enterprise adoption alongside the Office 365 suite.

The security controls include blocking unauthorised AI applications, preventing sensitive data sharing, and detecting potential intellectual property leaks through AI platforms.

https://www.cloudflare.com/

C3 Targets RPA Market with AI

C3 AI has launched what it describes as a next-generation robotic process automation platform that uses artificial intelligence agents to handle business and operational workflows without requiring coding expertise. The C3 AI Agentic Process Automation platform automates processes including order-to-cash, invoice processing, customer service, and supplier onboarding, as well as industrial operations such as equipment troubleshooting and production planning.

The company positions the product as an evolution beyond traditional RPA tools, which CEO Stephen Ehikian said follow "rigid instructions" requiring "substantial human intervention for escalations."

The new platform combines predetermined workflow steps with what C3 AI describes as AI reasoning capabilities.

"C3 AI Agentic Process Automation is different. It represents a step change in enterprise automation by combining the best of deterministic workflow steps with the dynamic reasoning capabilities of AI agents," said Nikhil Krishnan, C3 AI's chief technology officer for data science.

"This shift transforms automation from rigid scripts into intelligent systems that continuously deliver business value."

C3 AI claims all actions and steps are "fully transparent and auditable".

https://c3.ai

Dell AI Platform adds Unstructured Data Engine

Dell Technologies has updated its AI Data Platform with a new unstructured data engine developed through a partnership with search specialist Elastic, targeting enterprises struggling to harness unstructured data for artificial intelligence applications.

The updates address a critical enterprise challenge where only a fraction of rapidly-growing, unstructured data is currently usable for generative AI.

A key component of this new platform is an unstructured data engine, developed in partnership with Elastic. This engine is designed to provide secure, realtime access to large-scale datasets for inferencing, analytics, and intelligent search. It leverages the

Elasticsearch vector database, providing advanced vector search, semantic retrieval, and hybrid keyword search capabilities, which are essential for powering AI applications.

“The key to unlocking AI’s full potential lies in breaking down silos and simplifying access to enterprise data,” said Arthur Lewis, president of the Infrastructure Solutions Group at Dell Technologies.

“Collaborating with industry leaders like NVIDIA and Elastic to advance the Dell AI Data Platform will help organizations accelerate innovation and scale AI with confidence”.

Dell claims the collaboration with Elastic will deliver advanced vector search and semantic retrieval capabilities through Elasticsearch vector database technology. The vendor says this addresses the need for continuous indexing and vector retrieval engines that convert content into embeddings for semantic search.

The company also announced new PowerEdge R7725 and R770 servers featuring NVIDIA RTX PRO 6000 Blackwell Server Edition GPUs. Dell claims the R7725 will be the first 2U server platform to integrate the NVIDIA AI Data Platform reference design.

According to Dell, the NVIDIA RTX PRO 6000 offers up to six times the token throughput for large language model inference compared to previous generation hardware.

The unstructured data engine within the Dell AI Data Platform works alongside other tools, such as a federated SQL engine for structured data and a processing engine for large-scale data transformation.

This integrated approach is intended to transform massive datasets into reliable, high-quality, real-time intelligence for generative AI applications.

According to Ken Exner, Chief Product Officer at Elastic, “Fast, accurate, and context-aware access to unstructured data is key to scaling enterprise AI”.

He added that the partnership will bring “vector search and hybrid retrieval to a turnkey architecture, enabling natural language search, real-time inferencing, and intelligent asset discovery across massive datasets”.

The unstructured data engine and new server configurations will be available later this year, though Dell has not provided specific pricing or regional availability details.

https://www.delltechnologies.com

Hyland Unveils Enterprise Context Engine and Mesh

Content management vendor Hyland has unveiled two AI-powered technologies designed to automate enterprise workflows and decision-making across industries including healthcare, banking and government, insurance, government, and higher education.

Announcing its Enterprise Context Engine and Enterprise Agent Mesh, the company claimed the technologies represent what it calls “industry-first” solutions for linking organisational content, processes and applications.

The Enterprise Context Engine creates what Hyland describes as a unified view of enterprise operations by integrating data across systems including ERP, CRM and electronic health records.

The Enterprise Agent Mesh deploys multiple AI agents tailored to specific industry workflows.

The technologies build on Hyland’s existing Content Innovation Cloud platform, which the company positions as an alternative to rebuilding enterprise systems for AI implementation.

Hyland also announced a strategic collaboration agreement with Amazon Web Services to accelerate development of AI-powered document processing solutions for regulated industries.

The agreement expands Hyland’s existing technology partnership with AWS through enhanced cloud integration, joint go-to-market initiatives, and co-development of tools for managing unstructured data across healthcare, financial services, government, and insurance sectors.

The collaboration centres on Hyland’s Content Innovation Cloud platform, which the company says transforms unstructured content into actionable data using artificial intelligence and automation technologies.

Hyland claims the Content Innovation Cloud is “purpose-built for regulated industries” and features capabilities including semantic AI that “understands content with human-like reasoning.”

The strategic agreement represents a deepening of Hyland’s cloud-first strategy.

https://www.hyland.com

Korean AI Tool Tackles Workflow Automation

Seoul-based Infofla has launched Version 2 of its Selto automation platform, targeting persistent challenges with traditional robotic process automation that fails when user interfaces change unexpectedly.

The company claims its “VLAgent” engine combines large language models with visual recognition to enable automated workflows that adapt to interface modifications, pop-ups and layout changes that typically break conventional RPA deployments.

Traditional RPA implementations face significant challenges when applications update their user interfaces, with bots failing to complete tasks when buttons move or interfaces refresh. Industry analysts report that up to 50% of initial RPA projects fail, often due to the technology’s inability to adapt to changing environments.

Selto V2 introduces visual confirmation capabilities that allow users to see what the system has learned from screens, according to Infofla CEO In-mook Choi. The platform can now process external data sources and create conditional workflows based on user types or interface conditions.

The South Korean Ministry of the Interior and Safety has deployed Selto. The company reports it is conducting proof-of-concept trials with financial and insurance sector clients.

The platform’s ability to handle “unpredictable digital environments” represents a significant technical challenge that traditional RPA vendors are also addressing through AI integration. Major RPA providers are incorporating generative AI and machine learning to improve bot resilience and decision-making capabilities.

The company offers desktop installation allowing individual users to train automation workflows, potentially appealing to organisations seeking to democratise automation beyond IT departments.

https://www.infofla.com/en/home

API Tool Targets PDF Workflows

Datalogics subsidiary pdfRest has released a Sign PDF API Tool that applies cryptographic digital signatures to PDF documents through its REST API service. The US company’s new tool targets developers integrating PDF processing into applications and platforms.

The tool uses cryptographic processes rather than image overlays to create tamper-proof signatures. It supports industry-standard digital certificates including PFX format.

The company claims the signatures comply with international regulations including the US ESIGN Act and EU eIDAS regulation.

The API addresses growing demand for automated document signing workflows. Recent surveys indicate approximately three-quarters of organisations still use a mix of paper-based and digital document workflows.

Key applications include contract signing, internal approvals, and batch processing of documents such as invoices and certificates.

The API provides cryptographic timestamps and identity verification for audit trails.

Developers can test the tool through pdfRest’s API Lab browser interface, which generates code for integration.

The service competes with established providers including DocuSign and Adobe Sign.

https://pdfrest.com

5AI Agents Target Data Governance

Reltio has unveiled AgentFlow, an artificial intelligence platform designed to automate data governance and business operations through autonomous agents.

The cloud-based platform builds on Reltio’s existing Data Cloud infrastructure to deliver what the company describes as purpose-built agents for enterprise data tasks. These include resolving data matches, enriching attribution details, identifying quality anomalies, and validating compliance data.

AgentFlow agents operate with conversational interfaces and can orchestrate multiple sub-agents to handle complex workflows. The platform integrates with various large language models and includes role-based access controls for different user types, from data stewards to business analysts.

Early access customers include Radisson Hotel Group and Eaton Corporation, who are developing agents for match resolution, hierarchy management, and data quality enhancement. Global consulting partners Cognizant, ZS, and TCS are supporting client implementations.

“Enterprises urgently need agentic AI applications that go beyond experimentation,” said Manish Sood, Reltio CEO and founder. “Generic solutions can’t deliver real outcomes without the context of enterprise knowledge.”

The platform includes an MCP (Model Context Protocol) Server that allows enterprises to integrate Reltio agents with custom-built or third-party agents while maintaining data governance standards.

The MCP Server is available to customers now, while AgentFlow agents remain in early access with general availability planned for northern autumn 2025. Reltio has not disclosed pricing details.

https://www.reltio.com

Vanta Acquires Risk Monitoring Suite

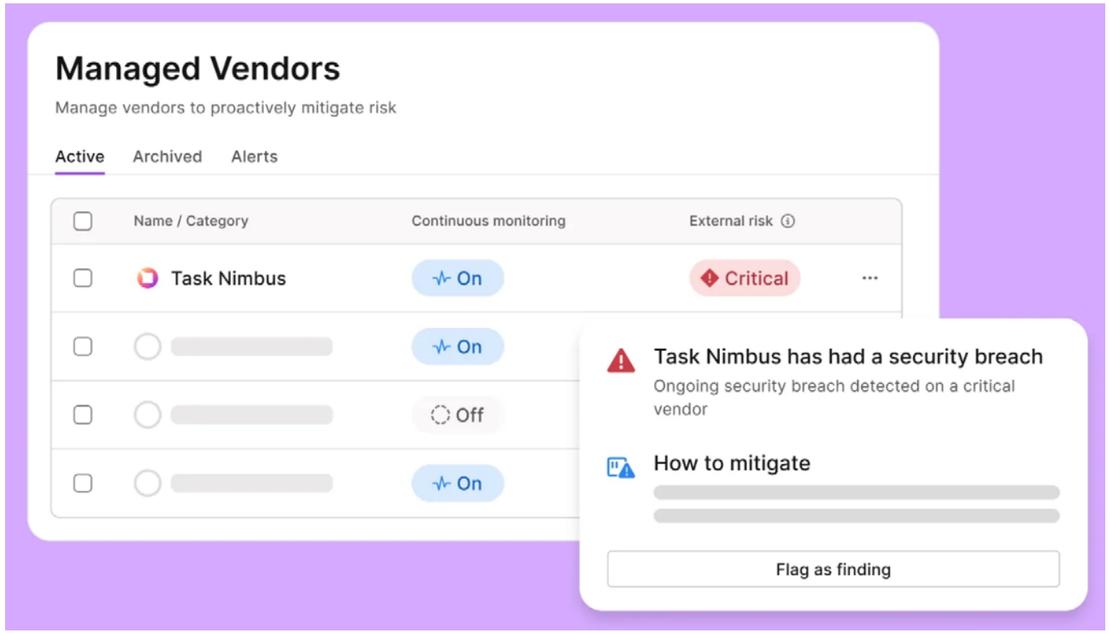

Vanta has acquired Riskey, a specialist in real-time third- and fourth-party risk monitoring, in a move to transform how organizations manage vendor cybersecurity risks. The acquisition integrates Riskey's technology into Vanta's Vendor Risk Management platform, replacing traditional static point-in-time assessments with continuous, AI-driven risk intelligence.

The deal comes as supply chain cyberattacks have surged dramatically, with supply chain attacks increasing by 431% between 2021 and 2023, and third-party involvement in breaches doubling to 30 percent according to recent industry reports. Supply chain attacks rose sharply in April-May 2025, hitting IT and telecom sectors hardest, demonstrating the urgent need for enhanced vendor monitoring capabilities.

The timing reflects growing pressure on organizations to better manage vendor relationships amid escalating cyber threats. According to Gartner research cited in the announcement, 45% of organizations have experienced increased business interruptions due to third-party cybersecurity incidents, while IT teams spend over six hours weekly reviewing vendor risk.

"The traditional model of vendor risk reviews—annual questionnaires and lagging scores—no longer meets the pace or scale of today's threat landscape," said Jeremy Epling, Chief Product Officer at Vanta. "By integrating Riskey and Vanta, we've unlocked continuous vendor risk monitoring which lets customers identify threats proactively and take action immediately to protect company assets."

Recent high-profile incidents have underscored the vulnerabilities in supply chains. The Change Healthcare cyberattack disrupted medical billing nationwide, delaying patient care and causing widespread financial strain, while attacks on Australian IT and telecom solutions companies exposed licensing files, hashed credentials, and critical infrastructure data.

Riskey's technology monitors for vulnerabilities, breaches, misconfigurations, leaked credentials and subprocessors across third- and fourth-party relationships. Its AI scoring model categorizes findings and provides actionable context to security teams.

"Customers are drowning in vendor data with no clear signal on what's relevant or actionable," said Koren Molcho, CEO and co-founder of Riskey. "Combining Riskey's monitoring technology into Vanta's VRM offering is an absolute game changer."

The integrated platform will enable customers to run comprehensive vendor assessments through both first-party security reviews and third- and fourth-party monitoring, streamline risk management with automated artifact collection, and receive real-time alerts with contextual guidance for mitigation.

https://www.vanta.com

Retab Raises $US3.5M for Document AI

San Francisco-based startup Retab has secured $US3.5 million in pre-seed funding to develop its document processing platform that extracts structured data from PDFs and scanned documents using artificial intelligence.

Retab's platform allows developers to define data schemas while the system handles document processing through automated model selection and prompt engineering. The company claims its approach can reduce costs compared to alternative solutions, though specific performance metrics could not be independently verified.

The startup was founded by engineers Louis de Benoist (CEO), Sacha Ichbiah, and Victor Plaisance, who previously built document automation tools for logistics workflows. The 10-employee company emerged from stealth mode with the funding announcement.

Retab positions itself as a developer-focused platform that works with existing large language models from OpenAI, Google, and Anthropic rather than developing proprietary models. The platform features what the company describes as "self-optimising schemas" and multi-model consensus mechanisms for accuracy validation.

The company claims its system can achieve "up to 100x" cost reductions compared to competitors and cites unnamed customers in trucking and financial services achieving high accuracy rates.

Retab plans to expand its platform to process web content and develop integrations with additional automation platforms. The company envisions serving as middleware between unstructured data and AI agents requiring structured information.

The startup aims to become foundational infrastructure for what it terms "vertical AI" applications across industries processing document-heavy workflows.

https://retab.ai

Risk compliance platform adds AI summarisation

Financial crime screening platform Sigma360 has launched Entity Summary, a generative AI feature that automatically creates risk profiles from multiple data sources.

The cloud-based tool uses large language models to consolidate information including KYC data, watchlist details, corporate registry information and adverse media intelligence into standardised summaries for compliance teams.

Sigma360 claims to be the first screening platform to offer this capability concurrently.

The feature addresses a common compliance challenge where analysts manually review and summarise fragmented risk data across multiple systems for due diligence reports. This process typically requires significant time investment and creates potential for inconsistency in reporting.

Entity Summary promises:

■ Accelerated Due Diligence: Immediate access to holistic entity summaries, eliminating time-intensive manual investigations across multiple systems and data sources.

■ Automated Case Reporting: Creates copyable report content that compliance teams can readily organize into professional due diligence reports.

■ Greater Consistency: Standardizes summary outputs in a concise and structured format, lowering chances of reporting errors and saving proofreading time.

Sigma360 is currently offering Entity Summary as a feature add-on to the company's existing flagship screening platform. Users can generate summaries instantly as part of their established compliance workflows.

This development fits broader industry trends toward automation in compliance operations, as organisations seek to reduce manual workloads while maintaining regulatory requirements.

https://www.sigma360.com/genai-solutions

Partnership Targets AP Compliance

Sovos and Tungsten Automation have formalised a strategic partnership to integrate accounts payable automation with tax compliance capabilities, responding to increasing global regulatory complexity. The collaboration combines Tungsten's AP automation platform with Sovos' tax compliance cloud to create what the companies claim is a comprehensive solution for invoice processing and global tax compliance.

The partnership builds on an existing relationship where Sovos previously served as an embedded tax compliance component within Tungsten's platform. The new agreement formalises this arrangement following what the companies describe as "direct customer feedback and operational success."

The integrated solution targets organisations facing new global tax mandates including the EU's VAT in the Digital Age (ViDA) initiative and Continuous Transaction Controls (CTC). While these are European requirements, similar digital compliance trends are emerging globally as tax authorities modernise their systems and reporting requirements.

In Australia, the Australian Taxation Office has been implementing realtime reporting requirements and digital service obligations, creating parallel compliance challenges for organisations with international operations.

The combined platform provides AI-powered optical character recognition for invoice capture, automated matching against purchase orders, realtime tax determination, and integration with ERP systems including SAP.

It provides global tax determination and ensures real-time compliance with e-invoicing regulations

The solution is positioned as source system-agnostic, connecting to existing infrastructure without requiring system replacement.

https://www.tungstenautomation.com/

Sydney AI startup's no-code automation

A Sydney-based artificial intelligence startup claims to bridge the gap between complex AI tools and practical business automation through a platform that requires no coding expertise.

Mazaal AI, founded in 2023 by former McKinsey QuantumBlack consultant Enod Bataa, has developed what it calls an "AI builder" that allows users to create automated workflows using plain English descriptions.

The company says it has backing from NVIDIA and support from the New South Wales government.

The platform uses what Mazaal AI terms "agentic" artificial intelligence - digital assistants designed to make context-aware decisions rather than simply executing predefined tasks.

The system employs Retrieval-Augmented Generation technology, combining large language models with company databases and documents.

"Businesses were drowning in repetitive work, but the tools to automate it were either too complex or too expensive," Bataa said, describing his motivation for founding the company.

The platform maintains a human-in-the-loop design, where AI agents can process information and draft responses but require human approval for key decisions.

For instance, when handling customer discount requests outside standard policy, the system might check inventory and order history before asking a manager for approval.

According to the company, the technology is being used by small retailers for inventory management, service providers for client onboarding, and larger enterprises for customer support functions.

https://mazaal.ai

Denodo AI tackles Governance

Data management vendor Denodo has announced the availability of DeepQuery in Denodo Platform 9.3. DeepQuery is a Deep Research capability that extends beyond fact-based retrieval to tackle complex, analytical questions with detailed, fully explained reasoning.

The platform update introduces dynamic access controls for realtime policy lookups, automated business context generation, and schema-resilient materialised views. These features target organisations struggling with AI project implementation amid changing compliance requirements and distributed data landscapes.

DeepQuery, pre-announced in July, is now generally available on GitHub under Apache licence. It enables multi-step analytical queries across enterprise systems. The tool promises to deliver insights in minutes that would typically require days of analyst work, according to the company.

Denodo cites an MIT study claiming 95% of AI projects fail to deliver positive return on investment, attributing this to inadequate data infrastructure supporting operational AI workloads. The company positions its platform as addressing bidirectional, real-time data access requirements for AI agents and chatbots.

Key Platform 9.3 enhancements include materialised views that adapt to schema changes, supporting rapid AI application development. The system now auto-generates metadata embeddings in vector databases and provides write-back capabilities to Iceberg tables through Databricks Unity.

DeepQuery represents Denodo's move toward open-source accessibility for AI developers. The tool leverages the company's existing semantic layer and AI SDK to orchestrate complex data retrieval and reasoning workflows.

The system maintains enterprise governance policies while enabling AI-driven queries, potentially addressing concerns about AI applications accessing sensitive data without proper oversight.

https://www.denodo.com

Cloud Document Service GenAI Boost

Fujifilm Business Innovation has introduced Generative AI capabilities to its IWpro cloud document processing service, targeting small and medium businesses struggling with non-standard form digitisation. The new Intelligent Data Capture Option uses AI-enhanced optical character recognition to extract data from business documents with varying layouts without requiring manual configuration.

The FUJIFILM IWpro Intelligent Data Capture Option utilizes generative AI to efficiently digitize data from quotations, purchase orders, delivery notes, and invoices used in business-to-business transactions. Its AI-enhanced OCR technology accurately extracts information appearing in line-items tables - despite wide variations in document formats - and converts it into structured digital data.

Structured data output is provided in formats compatible with downstream systems such as sales or accounting platform, enabling smooth processing of subsequent tasks including order registration, payments, and journal entries. The service eliminates the need for customer-side prompt configuration, reducing setup effort while helping streamline document processing operations.

Previously, extracting necessary data from "ledger sheets" in FUJIFILM IWpro would require either specifying reading positions (coordinates) or pulling information located around predefined keywords. While effective for standardized form, this approach was challenging for non-standardized, company-specific layouts.

It also limits the items that could be reliably captured – particularly for critical order information such as item names, quantities, and amounts within line-items tables that can vary in length and layouts. This often resulted in manual data entry.

To address this, FUJIFILM Business Innovation has enabled robust reading of line items tables in non-standard ledger sheets using OCR powered by generative AI.

The technology is able to structurally recognize text and tabular information while also interpreting the semantic meaning of each field.

For example, it will be able to associate "item-names" with its corresponding quantity and amount. As a result, the attributes and relationships of the data are preserved and output as structured digital information ready for use in subsequent systems.

FUJIFILM IWpro supports the full workflow from multifunction printer scanning and document capture to data extraction and output.

The service supports direct integration with Cybozu's Kintone platform and outputs data in formats for downstream business systems. Documents are processed through Fujifilm's cloud infrastructure.

https://www.fujifilm.com/fbau/en/products/au-software-products/fujifilm-iwpro

AI Data Filter Addresses Compliance Concerns

Enterprise data management specialist Komprise has launched its Intelligent AI Ingest system to help organisations safely filter unstructured data for artificial intelligence applications while maintaining compliance and security controls.

The company has announced general availability of the Smart Data Workflow ingestion engine, targeting enterprise concerns about data governance and sensitive information exposure in AI deployments.

The system addresses three critical challenges identified in Komprise's recent AI Data and Enterprise Risk survey: suboptimal AI outcomes from poor-quality data, high inferencing costs from irrelevant information, and security risks from bulk data ingestion that can expose sensitive information.

For compliance managers and CISOs, the platform includes built-in personally identifiable information detection and custom sensitive data classification to reduce compliance violations.

The system maintains automated audit trails documenting data lineage for regulatory reporting requirements.

Kumar K. Goswami, Komprise CEO, said the solution helps organisations "untangle the mess of unstructured data to gain the greatest competitive advantage with AI" while addressing customer concerns about efficiently moving appropriate data to AI systems.

The platform delivers a metadata-rich global file index providing enterprise-wide visibility of file data through simple queries. Komprise claims 2X performance improvement over major cloud provider data transfer tools through purpose-built transfer engines and massively parallel architecture.

The elastic grid architecture processes multiple network interfaces, share engines, and thread pools simultaneously, enabling rapid indexing across billions of files while moving large data volumes to various AI tools and services.

Research firm Gartner notes that modern Data Storage Management Services solutions are "foundational to business analytics and generative AI initiatives, helping enterprises unlock the full value of their data by making it more discoverable, contextualised and actionable."

https://www.komprise.com/

Morae Bolsters Records Management

Houston-based Morae Global Corporation has acquired Gimmal, a leading information governance software platform.

The acquisition integrates Gimmal's end-to-end information governance software with Morae's existing legal technology solutions, creating what executives describe as a comprehensive platform for managing the complete information lifecycle - from document creation and classification to retention and disposal.

Founded in 2002, Gimmal software addresses critical compliance and operational needs that have become increasingly important as organizations grapple with growing data volumes and evolving regulatory requirements.

"Our clients want complete, comprehensive solutions that address increasingly complex legal, regulatory, and operational demands," said Shahzad Bashir, Morae's Chairman and CEO.

"The combination of Gimmal's software with Morae's MorAI technology, services and consultancy will deliver an unmatched breadth and depth of offerings in our field."

The deal represents Morae's sixth acquisition since private equity firm Lateral Investment Management invested in the company in 2019, highlighting an aggressive growth strategy focused on building a comprehensive legal technology platform.

Craig Carpenter, CEO of Gimmal, emphasized the strategic value of joining forces with Morae's global reach and established client base.

"Plugging the leading information governance software into the leading, global legal technology company will deliver myriad new benefits and synergies for our global network of corporate legal departments and law firms," he said.

Financial terms of the acquisition were not disclosed.

https://www.morae.com/

NE2NE tackles complex PDF data extraction

NE2NE has released an AI-assisted tool, PDFFlex, designed to extract data from complex PDF files, targeting businesses struggling with manual data entry processes.

NE2NE claims the tool can reduce data extraction times from hours to minutes for documents with embedded tables. The company targets HR departments processing payroll registers and law firms conducting wage audits.

The tool extracts data from PDFs and converts it into Excel spreadsheets, XML or JSON formats. PDFFlex includes AI validation to check extracted data against original documents and sends alerts for human review when needed.

PDF data extraction remains a significant challenge for businesses managing compliance processes and digital workflows. Many organisations still rely on manual data entry from complex documents, creating bottlenecks in automated systems.

NE2NE, founded in 2021, describes itself as "the world's only fully agnostic data integration platform". The company focuses on small-to-midsize businesses seeking automation solutions.

The launch addresses growing demand for automated document processing as organisations digitise operations. However, specific technical details about the AI tools used and security compliance standards were not disclosed.

Company founder Steven Pappadakes said PDFFlex "deepens our product suite to offer a more comprehensive way for small-to-midsize companies to bring all their data integrations under one roof."

https://ne2ne.com

Stop Email Attacks on AI Assistants

Cybersecurity company Proofpoint has announced new protections against email-based attacks targeting AI assistants including Microsoft Copilot and Google Gemini, as organisations increasingly integrate artificial intelligence agents into workplace workflows.

The firm revealed four security innovations designed for what it terms "the agentic workspace" – environments where employees collaborate alongside AI agents and assistants.

Attackers are embedding malicious prompts in emails to manipulate AI assistants, using prompt injections to provide false information to users, confuse AI-based defences, and steal sensitive data, according to Proofpoint's announcement.

The company's Prime Threat Protection solution will block these exploits before they reach inboxes, enabling staff and AI agents to trust workplace interactions.

Proofpoint's Data Security Complete solution addresses growing data risks in AI-enabled workplaces by providing discovery, classification and control across endpoints, email, web and cloud platforms.

The solution includes Autonomous Custom Classifiers for dynamic data classification and creates consolidated risk maps showing cross-channel data lineage and potential exfiltration risks.

A companion AI Data Governance capability enables organisations to discover both authorised and unauthorised AI usage while applying policies to prevent data exfiltration and privacy violations.

The company's Secure Agent Gateway, built using Model Context Protocol, monitors and controls how customer-deployed AI agents access organisational data.

The gateway enforces data usage policies and can block or redact sensitive information before sharing between agents or with human colleagues.

Proofpoint Satori Agents operate within the company's security platform to handle data loss prevention alerts, recommend phishing simulations, and resolve user-reported email threats automatically.

"The agentic workspace is here and one of the most profound shifts in terms of how work gets done," said Proofpoint CEO Sumit Dhawan.

"Protecting the agentic workspace is the next evolution of human-centric security, extending beyond people to safeguard AI agents and the points where they collaborate and share data."

The AI exploit detection capability for email is expected in Q4 2025, while Data Security Complete is available from Q3 2025. The Secure Agent Gateway and Satori Agents enter phased availability from 2026.

http://www.proofpoint.com/

CData targets AI governance

Software vendor CData has released Connect AI, a managed platform connecting AI applications to more than 300 enterprise data sources while maintaining existing security protocols.

The company claims Connect AI addresses a critical challenge facing organisations deploying AI: accessing enterprise data while maintaining governance controls.

Connect AI inherits user permissions and authentication directly from source systems, with all data access logged under the authenticated user or agent's identity. Additional AI-specific controls can be layered within the platform, according to the company.

CData referenced MIT research indicating 95% of AI pilots fail to deliver measurable business impact, primarily due to data access and governance challenges.

The platform tackles two core issues: preserving contextual relationships within data that AI systems need for decision-making, and maintaining security protocols established in source systems.

"Enterprises that want to safely and effectively put their business data to work with AI need real-time access combined with semantic understanding. AI needs to comprehend what data means, not just where it lives," said Manish Patel, CData's Chief Product Officer.

Connect AI can be deployed in the cloud or embedded within software products using point-and-click configuration. Independent software vendors can white-label the offering within their products.

The platform uses connectivity technology already embedded by companies including Palantir, SAP, Salesforce Data Cloud and Google Cloud, repositioned for AI workloads with rea-time integration capabilities.

"Organizations are struggling to scale AI because data is often siloed, inconsistent, or poorly governed, creating risk and inefficiency. Many AI initiatives stall as companies wrestle with integrating multiple data sources while maintaining compliance," said Stephen Catanzano, Senior Analyst, Data Management, Enterprise Strategy Group.

"Tools like CData's Connect AI are emerging in response to these widespread market challenges, reflecting the company's vision to streamline AI-ready data access across enterprises."

https://www.cdata.com/ai

Active Directory Ransomware Gap

Cohesity and Semperis have released Cohesity Identity Resilience, a new product targeting Active Directory protection amid escalating ransomware attacks on identity infrastructure.

The solution, now available for purchase through Cohesity, combines Active Directory hardening, immutable backups, rapid recovery capabilities, and forensic analysis tools. It supports both on-premises Active Directory and cloud-based Microsoft Entra ID environments.

The product launch addresses a critical vulnerability in enterprise IT infrastructure. Active Directory has become the primary target in ransomware campaigns, with nearly 90 per cent of attacks involving AD compromise, according to industry research.

The average cost of a ransomware attack reached $US9.36 million in 2024, yet only 27 per cent of organisations maintain dedicated recovery solutions for identity infrastructure.

Cohesity Identity Resilience provides several core capabilities. The system scans Active Directory environments for indicators of exposure and identifies attack paths to privileged assets. It offers immutable backups with cyber vaulting options and automated recovery workflows. Post-breach analysis tools assess Active Directory integrity before production restoration.

When Active Directory is compromised, attackers can gain unrestricted access to IT resources, often resulting in ransomware infection, data theft, and prolonged business disruption.

"Active Directory remains one of the most targeted assets in modern cyberattacks. Cohesity Identity Resilience, powered by Semperis, enables organizations to strengthen their resilience, ensuring they can withstand, respond to, and recover from identity-based threats with confidence,' said Gerry Sillars, Semperis, Vice President Asia Pacific & Japan.

https://www.cohesity.com

AI Content Platform for Banks

Fiserv has announced Content Next, a cloud-based content management solution developed with OpenText that targets operational efficiency and compliance challenges in banking.

The multi-tenant software-as-a-service platform, built on OpenText's Core Content Management infrastructure, aims to automate document-intensive processes including loan processing, compliance reviews and customer onboarding.

The solution addresses automation priorities identified in KPMG's 2025 Banking Survey, where automated systems have reportedly reduced manual processing time by up to 60 per cent in critical banking operations. Content Next incorporates natural-language search, document classification and automated processing capabilities. The platform enables financial institution administrators to manage user permissions and configure access controls without IT support, according to the vendors.

"By integrating AI into operational workflows, institutions can speed up processes, and free up time and resources to focus on serving customers and members," said Whitney Russell, President of Digital and Financial Solutions at Fiserv.

Sandy Ono, EVP and Chief Marketing Officer at OpenText, said the partnership combines "OpenText's capabilities and Fiserv's extensive experience meeting the operating needs of financial institutions."

The platform includes role-based workspaces for compliance, loan underwriting, risk management and customer service teams. Built-in retention policies, version control and audit trails address governance requirements, while native integration with Microsoft 365 and Google Workspace supports existing technology environments.

Fiserv claims institutions using automated content workflows have reported 30-50 per cent faster document turnaround times and 25-40 per cent reductions in operational costs.

https://www.fiserv.com