idm.
information & data manager
FEBRUARY-MARCH 2026

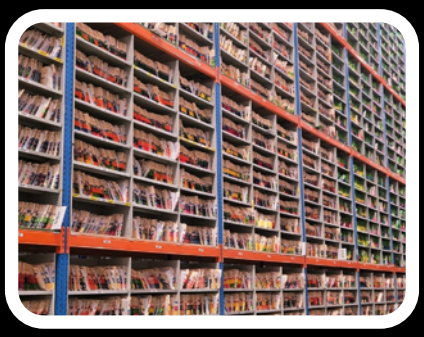

Digital Health & the Persistent Paper Trap

Cyber security's role in digital transformation

Does Your Business Need ISO 27001 Certification?

The Evolution of Data Governance in the Age of AI

AI Maturity is Top 2026 Tech Priority: Capgemini

Capgemini has identified 2026 as the year AI transitions from experimental deployments to enterprise-wide implementations, with organisations needing to demonstrate measurable returns on AI investments.

The consulting firm's TechnoVision Top 5 Tech Trends report highlights that many AI pilots have failed to deliver expected outcomes. Pascal Brier, Chief Innovation Officer at Capgemini, says the issue stems from business approach rather than technology limitations.

"As we look ahead to 2026, AI moves beyond experimentation and enters a phase of maturity. The upcoming year will see AI become the backbone of enterprise architecture, reshape software lifecycle development, and redefine cloud consumption.

"At the same time, enterprise systems are undergoing a fundamental shift toward intelligent operations, while tech sovereignty emerges as a strategic priority, driving organizations to build resilient interdependence," said Brier.

The report identifies five technology trends reaching inflection points in 2026, with direct implications for CIOs, information managers and compliance professionals managing digital transformation initiatives.

Enterprise AI Architecture

Long-term AI value will require enterprise-wide implementations rather than isolated use cases. Organisations must establish data foundations and infrastructure while focusing on human-AI collaboration. The shift demands investment in data readiness and workforce reskilling. The report concludes 2026 will be the moment to move from proof-of-concept to proof-of-impact, ensuring AI drives measurable outcomes, trust, and collaboration at scale, whilst laying the foundations for larger-scale transformation to follow.

AI-Generated Software

AI is reshaping software development lifecycles across industries. Developers will specify outcomes while AI generates and maintains components, shortening delivery cycles. However, governance and oversight remain critical to prevent security gaps. The shift requires reskilling development workforces in systems thinking and AI orchestration.

Organizations will start rebuilding their applications and need to focus on reskilling their software development workforce in the near future. The new currency of expertise will instead lie in systems thinking, AI and agents orchestration, and managing complex, autonomous process and tool chains.

Cloud 3.0 Architecture

The analyst firm believes Cloud is entering its next evolution, a phase where hybrid, private, multi-cloud and sovereign architectures are no longer niche, but fundamental to how AI runs at scale, to the point it is becoming the operational backbone for AI and agentic workloads. Agentic systems require scalable, low-latency infrastructures with edge and cloud working as integrated fabric. Organisations must redesign architectures for performance, portability and strategic autonomy while managing increased complexity across diverse cloud environments.

Intelligent Operations

Enterprise systems are evolving from static records into adaptive operational engines. AI agents embedded in core processes will monitor activity, optimise execution and orchestrate workflows across finance, supply chain, HR and customer service. Automation shifts to human-AI co-steering where AI proposes and executes while humans supervise. Success depends on ensuring reliability and scalability of AI agents.

Technology Sovereignty

Tech sovereignty has moved from policy concept to strategic priority as nations and enterprises seek control over critical technologies.

Full autonomy does not exist, so organisations must focus on risk mitigation through diversified suppliers and sovereign alternatives. Securing business continuity becomes imperative through multi-clouds, regional AI models and open platforms.

The report notes that organisations need appropriate skills, agile governance and adaptive mindsets to operate confidently across diverse cloud environments while managing compliance requirements.

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

IDP Market Surges as AI Drives Automation

The global Intelligent Document Processing (IDP) market exceeded $US8 billion in 2024, growing 14.5% as generative artificial intelligence transforms how organisations automate compliance and data workflows, according to Infosource's annual State of the Global Intelligent Document Processing Market report.

The Swiss research firm projects the market will expand at 16% annually through 2029, driven by AI capabilities that eliminate template-based document processing and enable autonomous business process automation.

"IDP has evolved from a niche capability into a strategic foundation for intelligent automation," said Petra Beck, Senior Analyst at Infosource.

"GenAI and agentic orchestration are not just incremental improvements; they are redefining how organizations achieve efficiency, compliance, and resilience in a rapidly changing regulatory and business environment."

The technology shift enables template-free document extraction and multimodal input processing, allowing systems to handle diverse document types without pre-configuration.

This addresses a critical pain point for organisations managing complex regulatory submissions and compliance documentation across multiple formats.

The report concludes some organizations may delay investments in the near term to address AI skill gaps, refine deployment strategies, and adapt to evolving regulatory frameworks. However, adoption is expected to accelerate later in the forecast horizon.

Enterprises will expand GenAI-infused IDP deployments, apply domain-tuned language models and retrieval-augmented generation (RAG), and advance toward agentic orchestration once governance and ROI requirements are addressed.

Vendors embedding IDP within broader automation platforms are outperforming traditional capture-focused providers, the report found.

Agentic AI capabilities drive additional value through autonomous automation and adaptive optimisation that reduces manual intervention in document workflows.

Regulatory frameworks including the EU AI Act and expanding e-invoicing mandates are accelerating adoption globally. In the ANZ context, organisations preparing for mandatory e-invoicing adoption and enhanced data breach notification requirements under privacy legislation are investing in automated document processing capabilities.

However, some organisations may delay investments to address AI skills gaps and refine deployment strategies, Infosource warned. Governance, cost control, and return on investment emerge as critical success factors for GenAI implementations.

The report defines IDP as software and services that transform unstructured business documents into structured data for transactions, analytics, and compliance.

This capability becomes increasingly vital as organisations digitise paper-based processes while maintaining audit trails and regulatory compliance.

Data Readiness Gap Threatens AI Ambitions

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

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

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

IBM claims 81% of surveyed CDOs report their data strategy now integrates with technology roadmaps and infrastructure investments, up from 52% in 2023. However, only 26% express confidence in using unstructured data to deliver business value. The study notes 81% of respondents bring AI to data rather than centralising it.

While 92% of CDOs say they must focus on business outcomes to succeed, only one-third strongly agree they can clearly convey how data facilitates business results.

Just 29% have clear measures to determine the value of data-driven business outcomes, according to the study.

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

IBM reports 84% of surveyed CDOs say their unique data products have provided significant competitive advantages, with 78% citing proprietary data as a top strategic objective. The study found 80% of leaders have started developing diverse datasets to train AI agents. However, 79% admit being early in defining how to scale and govern them.

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

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

The research indicates 77% of leaders struggle to fill key data roles. Only 53% say recruitment efforts deliver needed skills and experience, down from 75% in 2024.

While 82% of CDOs say data is wasted without employee access, and 80% cite data democratisation as enabling faster organisational movement, fostering a data-driven culture remains a top strategic challenge.

Cyber security's role in digital transformation - why early engagement matters

By Leon Fouche

Digital transformation is reshaping industries, with organisations investing heavily in cloud platforms, artificial intelligence (AI), machine learning (ML), and data analytics to drive growth, resilience, and agility. However, cyber security is too often treated as a downstream safeguard rather than a strategic enabler.

According to a recent IDC report, sponsored by BDO, only 40 per cent of organisations integrate cyber security during the planning stage of digital initiatives. Most add security measures during execution or after implementation, significantly increasing the risk of costly and disruptive rework, project delays, and erosion of customer trust.

The gap between ambition and execution is visible in Australia, where governments and regulators are actively encouraging digital innovation across critical sectors. Yet, organisations face growing obligations under Australia's regulatory security requirements, including the Cyber Security Act, the Security of Critical Infrastructure Act (SOCI), Privacy Act reforms, and related data protection and critical infrastructure frameworks.

Embedding cyber security early in the digital transformation journey, and ensuring success, will require:

- Aligning cyber budgets with business strategy
- Refreshing cyber programmes to stay relevant
- Building cyber maturity for resilience.

Budget optimisation: Aligning spend with strategy

IDC's research highlights that budget is no longer the primary barrier - effectiveness of investment is. Even organisations with 'flexible' or 'readily available' budgets report an average of five incidents annually.

In Australia, this finding mirrors trends across both public and private sectors, where large cyber security investments have not always translated into fewer incidents or faster recoveries. The issue isn't always lack of funding, it's how that funding is applied.

What matters most is alignment of spend to strategic outcomes. Effective cyber security investment supports capabilities such as:

- 24x7 threat monitoring and response
- Automation in detection and response
- Proactive vulnerability and third-party risk management
- Integration of security into DevOps and cloud transformation programs.

For example, as Australian organisations pursue cloud migration under government digital strategies, it is essential that budgets extend beyond application redesign and migration into secure coding practices, regulatory risk assessments, and cloud security posture management.

Late stage bolt-ons almost always result in higher costs and reputational risks. To maximise impact, cyber security must be treated as a strategic partner, not a reactive fix. This need for strategic alignment naturally leads to the next consideration: how often organisations pause to reassess their cyber security approach.

IDC's findings show that boards increasingly demand proof of risk reduction, although many organisations lack process-level metrics to demonstrate genuine maturity. Boards don't just want compliance with frameworks, they want assurance that cyber risks are being actively reduced and that investment is delivering measurable outcomes.

To achieve this, organisations need to strengthen their understanding of risk and adopt security standards that are appropriate for their business context. Standards provide the structure to translate risks into actionable controls and measurable outcomes. Most Australian organisations do not rely on a single standard. Instead, many adopt a hybrid approach - aligning to ISO 27001 for certification, using the ACSC Essential Eight for baseline controls, and mapping to NIST CSF or sector frameworks such as AESCSF or APRA CPS 234 to meet regulator and industry obligations.

Cyber leaders must also regularly reassess their strategies to ensure they remain aligned with business priorities. Annual refreshes, outcome-based metrics, and cross-functional collaboration help teams stay relevant and effective. Reflection can also reveal legacy practices that dull progress. By shifting to agile, business-aligned approaches, cyber security teams can foster innovation and drive better outcomes. This practice strengthens collaboration between cyber and business units, breaks down silos, and builds trust.

Ongoing reassessment and alignment with regulatory change ensure cyber remains a business enabler. For example, state government agencies refreshing their digital roadmaps are increasingly embedding cyber strategy checkpoints into program stage-gates, ensuring controls remain relevant and adaptive to new risks.

Budget size doesn't guarantee security. IDC's survey reinforces that process maturity is the strongest predictor of resilience. Organisations that focus on 24x7 monitoring, advanced detection, and mature response processes achieve far greater resilience than those that spread budgets across point solutions. C

apabilities such as Extended Detection and Response (XDR), AI-driven analytics, and predictive modelling directly reduce incident frequency and accelerate recovery, proving that process maturity is the real driver of outcomes.

As cyber threats continue to evolve, future priorities must build on this foundation. Automation, endpoint protection for hybrid workforces, and employee awareness remain essential. At the same time, organisations are expanding focus to zero trust, disaster recovery, and supply chain resilience to address geopolitical and operational risks.

Emerging technologies such as Generative AI (GenAI) amplify the challenge: they offer powerful new capabilities but also fuel risks like phishing, data leakage, and governance gaps.

While many organisations are beginning to train staff and deploy AI-specific security tools, only a minority have embedded risk frameworks or governance processes. This highlights a clear maturity gap that must be closed before innovation can be embraced safely.

While awareness of risks is high, consistent delivery is lacking. Organisations that find a way to bridge the gap between strategy and execution hold a competitive advantage. To progress, cyber security must evolve from reactive control to intelligence-led resilience. This means embedding governance, enabling continuous monitoring for real-time visibility, and measuring process-level KPIs such as detection times, containment speed, and patching cadence. By making these practices core to operations, organisations move beyond compliance to a state of proactive, adaptive maturity - one that withstands today's threats and prepares them to adopt emerging technologies with confidence.

Cyber security must be embedded early in digital transformation

Treating cyber security as a strategic enabler from the planning stage reduces costly rework, delays, and reputational risks. Early integration ensures alignment with regulatory obligations and builds trust.

Align cyber budgets with strategic outcomes

Effective cyber investment isn't about spending more - it's about spending smart. Organisations must prioritise capabilities like threat monitoring, secure DevOps, and cloud security posture management to maximise ROI.

Cyber maturity drives resilience and innovation

Process maturity, not budget size, is the strongest predictor of cyber resilience. Embedding governance, automation, and real-time monitoring enables organisations to confidently adopt emerging technologies like GenAI.

Contact our cyber security team to discuss your options.

Leon Fouche is National Leader, Cyber Security, BDO Australia.

Originally published here.

Privacy risk is now a business risk, but investment lags

By Jamie Norton, Vice Chair, ISACA Board

Across Australia and New Zealand, privacy teams are being asked to shoulder one of the most complex risk environments organisations have ever faced. Rapid technological change, expanding regulatory obligations, the rise of artificial intelligence and an unrelenting wave of cyber threats have combined to make privacy a central business issue. Not just a compliance function tucked away in legal or IT.

Yet despite this growing responsibility, many organisations are quietly pulling back on the very resources needed to manage these risks effectively.

ISACA's newly released State of Privacy 2026 survey puts the numbers behind what many professionals are already feeling. Nearly two-thirds of privacy professionals in Oceania say their roles are more stressful today than they were five years ago. Technology change has become the leading source of pressure, followed closely by compliance challenges and resource shortages.

At the same time, budgets are heading in the wrong direction. Only eight per cent of Oceania respondents expect a privacy budget increase in the year ahead, while 60 per cent anticipate cuts.

This widening gap between expectations and investment should concern every business leader and board member. Privacy directly affects an organisation's reputation, customer trust, financial performance and governance standing. Data breaches carry immediate costs such as remediation, legal action or regulatory penalties, but the longer-term erosion of confidence can be far more damaging.

When privacy programs are under-resourced, the likelihood of failure increases. And the ISACA survey shows exactly where those failures tend to occur.

More than half of global respondents pointed to inadequate or poor training as a major cause of privacy breakdowns. Half cited the absence of privacy by design, which is focused on embedding privacy considerations into systems and processes from the outset. Nearly half pointed directly to breaches and data leakages.

I want to reinforce that these are not technical oversights. They are symptoms of organisations struggling to keep up with risk, using shrinking teams and limited funding. The workforce challenge is equally stark. Globally, the median size of privacy teams has dropped from eight staff to five in just one year. Both technical and legal expertise are in short supply, with over half of respondents identifying significant skills gaps, particularly in emerging technologies.

To cope, many organisations are retraining staff from other disciplines or relying more heavily on contractors and consultants. While career transitions into privacy can be a positive development, they also reflect how difficult it has become to build stable, specialised teams in a tightening labour market.

Meanwhile, confidence in privacy programs is slipping, particularly in our region. Only 26 per cent of Oceania respondents expressed strong confidence in their organisation's ability to protect sensitive data, well below the global average.

This matters because the risk landscape is only becoming more complex. AI adoption is accelerating across industries, bringing enormous opportunity but also introducing new data risks, transparency concerns and regulatory scrutiny. Against this backdrop, expecting smaller teams with fewer resources to manage enterprise-level risk is unrealistic.

Perhaps most concerning is the apparent retreat from foundational privacy practices. Fewer organisations report consistently applying privacy by design principles, and while security controls like encryption and data loss prevention remain widely used, even these are seeing slight declines.

On their own these shifts may appear minor, but over time they weaken an organisation's overall resilience.

Privacy risk is now everyone's responsibility across the organisation. It touches every business function from product development and marketing to HR, supply chains and customer engagement.

Boards and executives must quickly recognise that privacy capability is now as essential as financial controls, cyber security and operational resilience.

This starts with adequate funding, but it also requires clearer accountability. Privacy leadership needs a seat at the table where technology decisions are made. Risk assessments must extend beyond compliance to consider reputational, ethical and operational impacts. Training must be continuous and organisation-wide, not a once-a-year exercise.

Encouragingly, the survey also shows more organisations exploring the use of AI to support privacy functions, from automating compliance tasks to improving monitoring and risk detection. Used responsibly, these tools can help overstretched teams work more effectively. But they are not a substitute for skilled professionals and strong governance.

The message I hear from privacy professionals is straightforward - the job is getting harder, the risks are growing, and the resources are shrinking. That combination is not sustainable. Leaders who recognise this are already putting themselves in a stronger position. Organisations that invest in the right people, processes and technology, and embed privacy into everyday decision-making, will be far better placed to manage risk and maintain trust.

Jamie Norton is Vice Chair of the ISACA Board and Chief Information Security Officer at the Australian Securities and Investments Commission (ASIC). With more than 25 years of experience across government, commercial and international sectors, he specialises in cybersecurity, resilience and strategic risk management. Jamie has previously served as a Partner at McGrathNicol, advising executives and boards on navigating emerging technology risks; CISO at the Australian Taxation Office and has held leadership roles with NEC, Tenable, Check Point and the World Health Organization.

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

NZ Energy Firm Unifies Data for AI

New Zealand's largest electricity retailer Genesis Energy has selected the Databricks Data Intelligence Platform to centralise its data ecosystem and scale artificial intelligence applications across the organisation.

The Auckland-based energy company will unify disparate data sources on the platform as part of its broader Gen35 digital transformation strategy.

Genesis aims to deploy AI across asset management, trading operations, customer service and internal productivity tools.

The initiative builds on Genesis's AI Enablement Hub, established earlier in 2025 to standardise and govern enterprise AI deployment. The centralised hub addresses a common challenge for organisations scaling AI: maintaining consistent governance while enabling distributed innovation.

"By treating data as a strategic asset, we're driving real outcomes, accelerating the energy transition, and ensuring Genesis remains future-fit in a rapidly changing world," Genesis CEO Malcolm Johns said.

Chief Technology and Transformation Officer Ed Hyde said Databricks would accelerate the company's ability to embed data and AI into operations while upskilling staff.

Genesis will train more than 100 employees through the Databricks Data + AI Academy in coming months.

The upskilling programme targets employees across business units, reflecting the company's goal of democratising data access beyond IT departments.

The Databricks platform runs on an open-source foundation, allowing organisations to avoid vendor lock-in while building analytics and AI applications.

More than 20,000 organisations globally use the platform, including 60 per cent of Fortune 500 companies.

Mediaflux to Preserve Audiovisual History

The National Film and Sound Archive (NFSA), has selected Arcitecta's Mediaflux platform as its central Digital Asset Management System.

This move marks a significant step in NFSA's digital transformation journey to enhance access to the nation's cultural heritage.

The Mediaflux platform will enable NFSA to manage petabyte-scale collections, automate metadata enrichment and streamline digital workflows. It replaces legacy systems that struggled with the growing scale and complexity of digital assets.

"The NFSA's digital transformation isn't just about upgrading systems — it's about fundamentally rethinking how we manage, understand and share Australia's audiovisual history," said Dr. Keir Winesmith, Chief Digital Officer, NFSA.

"Mediaflux gives us the flexibility to structure and enrich our digital assets in ways that make them more discoverable, understandable and accessible.

"We're not just preserving history; we want to connect Australians with their audiovisual heritage and help them understand the many ways in which the past contains the present and the future."

NFSA selected Mediaflux following a rigorous evaluation, where the platform demonstrated unified data management capabilities and scalable, metadata-driven architecture. The implementation will proceed in multiple phases with initial focus on three key areas.

The first phase will deliver intelligent metadata management to improve searchability, next-generation collections management for streamlined curation, and a contributor portal for external deposits.

These improvements aim to better serve both internal workflows and public accessibility requirements.

Wasabi's cloud storage integration, managed by Mediaflux, will provide cost-effective scalability with no egress fees.

Dell PowerScale solutions will power high-speed digitisation and preservation of archive collections.

The project supports NFSA's broader strategy of enhancing discoverability and access while helping digitise at-risk materials across multiple National Cultural Institutions.

NSW Treasury eInvoicing tender

NSW Treasury has reappointed MessageXchange to deliver Peppol eInvoicing services following a competitive tender process.

The portfolio transferred from NSW Department of Customer Service to Treasury, which conducted the review. MessageXchange was originally appointed in 2020.

Peppol eInvoicing enables businesses to exchange invoices electronically between software systems, eliminating manual processing.

The technology aims to reduce costs and improve accuracy for government agencies receiving supplier invoices.

MessageXchange also provides eInvoicing services to the Australian Taxation Office, New Zealand's Inland Revenue and unnamed state governments.

The company states most small and medium businesses can already send eInvoices through major accounting software.

The announcement references ATO and Deloitte Access Economics research estimating $A28 billion in economic benefits from widespread eInvoicing adoption over 10 years. This figure represents potential economy-wide benefits rather than NSW-specific outcomes.

The release does not specify how many NSW agencies currently use the service, transaction volumes processed or quantified cost savings achieved.

https://www.messagexchange.com

Cybersecurity Chiefs Question Autonomous AI Safety

European cybersecurity authorities have issued new guidelines warning organisations against deploying fully autonomous artificial intelligence systems without human oversight, citing significant security risks that current technology cannot adequately address.

AI systems face three primary attack categories: Evasion Attacks, Poisoning Attacks and Privacy Attacks. Of particular concern is Indirect Prompt Injection, where attackers embed hidden instructions within text or data that AI models process without user awareness, potentially causing data leaks, incorrect decisions, or unauthorised actions.

The German Federal Office for Information Security and France's cybersecurity agency released joint design principles in August 2025 for securing Large Language Model systems using Zero Trust architecture. The 16-page framework targets the growing deployment of "agentic" AI systems that can operate independently across business processes.

"Blind trust in LLM systems is not advisable, and the fully autonomous operation of such systems without human oversight is not recommended," the agencies stated. "It is improbable that such agents can ensure meaningful and reliable safety guarantees."

Six Core Security Principles

The framework establishes six fundamental security measures:

■ **Authentication and Authorization** requires multi-factor authentication and role-based access controls, with particular emphasis on preventing LLM-based authentication systems.

■ **Input and Output Restrictions** mandate validation of all data flows, with gateways to detect malicious prompts and prohibit automatic preloading of external content.

■ **Sandboxing** isolates AI systems from unintended external interactions, including strict memory separation between user sessions and emergency shutdown capabilities.

■ **Monitoring, Reporting and Controlling** involves continuous threat detection, automated responses, and token limits to prevent resource abuse.

■ **Threat Intelligence** encompasses collecting and analysing emerging cyber threats, including regular red-teaming exercises and dynamic threat analysis.

■ **Awareness** requires comprehensive stakeholder education about AI risks, transparent decision-making processes, and regular security training programmes.

The guidelines specifically address vulnerabilities in Retrieval-Augmented Generation systems, where AI models access external databases, and warn against automatic execution of AI-generated system commands. "The user must be able to approve all system inputs of the application and actions of the agent," the agencies recommend.

Risk scenarios include data exfiltration through manipulated links, privilege escalation attacks, and supply chain compromises targeting AI system components. Critical mitigations include implementing least privilege access controls, comprehensive session isolation, and human-in-the-loop approval for sensitive operations.

The full report "Design Principles for LLM-based Systems with Zero Trust" is available for download here.

Leading AI Models Fail Accuracy Tests

A comprehensive evaluation of 37 major AI language models reveals significant weaknesses in factual accuracy that could pose compliance and operational risks for organisations deploying artificial intelligence tools. The study by Hong Kong University's Business School found that while leading models like GPT-5 and Claude 4 Opus performed best overall, all tested systems struggled with "factual hallucinations" – generating plausible but incorrect information that contradicts real-world facts.

Professor Jack Jiang, who led the research through the Artificial Intelligence Evaluation Laboratory, said hallucination control capability directly impacts the credibility of AI systems in professional settings including knowledge services, customer service and intelligent navigation.

The evaluation tested models on two types of hallucinations: factual errors that conflict with real-world information, and faithful errors where models fail to follow user instructions or produce content contradictory to input context.

Results showed GPT-5 variants achieved the highest overall scores of 86 and 84 respectively, followed closely by Claude 4 Opus models at 83 and 80. However, even top-performing models scored below 75 on factual accuracy tasks, indicating room for improvement in enterprise-critical applications.

For compliance and risk managers, the findings highlight potential vulnerabilities when deploying AI tools for document analysis, regulatory reporting or customer communications where factual accuracy is paramount.

The study found models generally excelled at following instructions precisely but were more prone to fabricating facts – a pattern that could mislead decision-makers relying on AI-generated insights for business-critical processes.

Chinese models including ByteDance's Doubao 1.5 Pro showed balanced performance but lagged behind international leaders, while reasoning-focused models performed better than general-purpose versions at avoiding hallucinations.

The research comes as organisations increasingly integrate AI capabilities into Microsoft 365 and other enterprise platforms, making hallucination control a critical consideration for digital transformation initiatives.

Information managers implementing AI workflows should establish validation processes and human oversight mechanisms to mitigate risks from factual inaccuracies, particularly in regulated industries where compliance failures carry significant penalties.

The full evaluation methodology tested models on information retrieval, misinformation identification and contradictory prompt scenarios to assess their ability to maintain factual consistency and contextual accuracy.

The full report is available here.

Microsoft Security Chief Warns of AI 'Double Agent'

Microsoft's Executive V-P of Security Charlie Bell has issued a stark warning about the security implications of AI agents in enterprise environments, urging organisations to implement robust governance frameworks to prevent AI from becoming "double agents" that undermine cybersecurity efforts.

In a blog post, Bell emphasised that while AI promises unprecedented productivity and innovation, it also introduces unique security risks as organisations rapidly deploy AI agents across their operations.

"AI isn't just another chapter - it's a plot twist that changes everything. The opportunities are huge, but so are the risks," Bell wrote.

Drawing a parallel to Star Trek characters, Bell compared the dual nature of AI to the android Data and his evil twin Lore, highlighting how AI agents can either strengthen or compromise security postures.

The warning comes as IDC research predicts there will be 1.3 billion AI agents in circulation by 2028, creating an urgent need for enhanced security measures.

Bell outlined three key principles for managing AI security risks: recognising the new attack landscape, practicing "Agentic Zero Trust," and fostering a culture of secure innovation.

"Unlike traditional software, AI agents are even more dynamic, adaptive and likely to operate autonomously. This creates unique risks," Bell explained.

The Microsoft security chief advocated for applying Zero Trust principles to AI deployments through "Containment" and "Alignment" - concepts he attributes to discussions with Mustafa Suleyman, Executive Vice President and CEO of Microsoft AI.

"Containment simply means we do not blindly trust our AI Agents, and we significantly box every aspect of what they do," Bell noted, adding that organisations must never let "any agent's access privileges exceed its role and purpose."

Bell emphasised the importance of proper identity management for AI systems, stating that "every agent must have an identity" with clear accountable ownership within the organisation.

To combat emerging threats, Bell recommended several practical steps, including assigning each AI agent an ID and owner, documenting their intent and scope, monitoring their actions, and keeping them in secure, sanctioned environments.

Microsoft has been developing solutions to address these challenges, including Microsoft Entra Agent ID, which helps customers assign unique identities to agents created in Microsoft Copilot Studio and Azure AI Foundry.

Global DDoS attack surge threatens critical infrastructure

Distributed Denial of Service (DDoS) attacks jumped globally by 40% year-over-year in the third quarter of 2025, with security provider Cloudflare blocking 8.3 million attacks, averaging 3,780 attacks per hour - according to the company's latest quarterly threat report.

The report's release follows Cloudflare's own significant network outage on 18 November 2025. The company claims the three-hour disruption was not caused by a DDoS attack but by an internal database configuration error.

The incident initially resembled an attack from the Aisuru botnet, causing teams to investigate potential DDoS activity before identifying the internal issue.

"An outage like today is unacceptable," Cloudflare stated in its incident report. The outage affected core CDN services, Turnstile, Workers KV, dashboard access and email security functions and underscores vulnerabilities even major security providers face.

Cloudflares's Q3 threat report reveals an escalating attack landscape driven primarily by the Aisuru botnet, estimated to control 1-4 million infected devices globally.

Aisuru launched 1,304 hyper-volumetric attacks during Q3 2025, representing a 54% increase from the previous quarter. These attacks routinely exceeded 1 terabit per second (Tbps) and 1 billion packets per second (Bpps).

The botnet achieved a record-breaking 29.7Tbps attack using UDP carpet-bombing techniques, bombarding an average of 15,000 destination ports per second. Cloudflare's autonomous mitigation systems detected and blocked the attack without human intervention.

"If Aisuru's attack traffic can disrupt parts of the U.S. Internet infrastructure when said ISPs were not even the target of the attack, imagine what it can do when it's directly aimed at unprotected or insufficiently protected ISPs, critical infrastructure, healthcare services, emergency services, and military systems," the report stated.

Security researcher Brian Krebs reported the botnet caused "widespread collateral Internet disruption" in the United States when attack traffic routed through Internet service providers. Portions of the Aisuru botnet are available for hire, enabling attackers to launch nation-scale disruptions for several hundred to several thousand U.S. dollars.

Network-layer attacks surged 95% year-over-year to 5.9 million incidents, accounting for 71% of all DDoS attacks in Q3. HTTP DDoS attacks decreased 17% year-over-year to 2.4 million attacks.

The threat landscape highlights vulnerabilities in legacy DDoS protection systems. Most attacks (89% at network-layer, 71% at HTTP layer) conclude within 10 minutes - too fast for human response or on-demand mitigation services to activate effectively.

Short-lived attacks create particular risks for organisations managing compliance and digital transformation initiatives. Disruption extends beyond the attack duration, requiring complex recovery processes including system restoration, distributed data consistency checks and secure service rebuilding.

Indonesia maintained its position as the largest source of DDoS attacks globally, a ranking it has held since Q3 2024. HTTP DDoS attack requests originating from Indonesia have increased 31,900% over five years.

The automotive industry experienced the largest surge, jumping 62 positions to become the sixth most-attacked industry globally. This coincided with escalating EU-China trade tensions over electric vehicle tariffs and rare-earth mineral exports.

The mining, minerals and metals sector surged 24 positions, while cybersecurity companies climbed 17 spots as attacks intensified. DDoS traffic against AI companies spiked 347% month-over-month in September 2025 as regulatory scrutiny of artificial intelligence increased.

Information technology and services topped the most-attacked industries list, followed by telecommunications and gambling sectors. Geopolitical events correlated directly with attack patterns, including protests in the Maldives, France and Belgium that coincided with significant DDoS activity increases.

UDP attacks increased 231% quarter-over-quarter, driven largely by Aisuru activity, making it the primary network-layer attack vector. DNS floods ranked second, followed by SYN floods and ICMP floods, collectively accounting for over half of network-layer attacks.

The Mirai botnet, despite first appearing nearly a decade ago, still launches almost 2% of network-layer DDoS attacks. Nearly 70% of HTTP DDoS attacks originated from botnets already catalogued by Cloudflare's threat intelligence systems.

Cloudflare has mitigated 36.2 million DDoS attacks in 2025 through the end of Q3, representing 170% of the total attacks blocked throughout 2024. The company provides unmetered DDoS protection to all customers regardless of attack size, duration or frequency.

The full report is available here.

Balancing Compliance, Security, and AI

The Top Question for 2026 for Records Managers

By Brandon Voight

As government agencies and enterprises across Australia and New Zealand accelerate digital transformation, one question dominates the minds of records managers: 'How do we ensure compliance and security in an AI-driven, cloud-first environment while meeting retention and disposal requirements?'

This challenge reflects a convergence of factors - rapid cloud adoption, the rise of generative AI, escalating cyber threats, and increasingly complex regulatory frameworks.

For records managers, the stakes have never been higher. Failure to adapt could mean compliance breaches, reputational damage, and significant financial penalties.

The Compliance Landscape: A Moving Target

In Australia, compliance obligations are evolving under multiple frameworks. The Public Records Act 2023 in Queensland introduced new governance requirements, with mandatory standards expected after mid-2025. Agencies must now demonstrate robust information governance while preparing for stricter retention and disposal rules.

Similarly, State Records NSW has announced biennial monitoring exercises starting in 2026, requiring agencies to self-assess maturity using the Records Management Assessment Tool (RMAT).

These assessments will be verified by State Records NSW, signaling a shift toward greater accountability and transparency in recordkeeping practices.

Across the Tasman, New Zealand faces its own compliance pressures.

The latest Chief Archivist's Annual Report highlights persistent gaps in adherence to the Public Records Act 2005, compounded by shrinking information governance teams - down more than 15% over the past two years.

Agencies are struggling to dispose of digital records systematically, creating long-term risk exposure.

Cloud and AI: Opportunity Meets Risk

The Australian Government's Data and Digital Government Strategy 2025 Implementation Plan underscores the push toward cloud-first services and AI-enabled decision-making.

The plan prioritizes artificial intelligence, data integration, and cyber resilience to deliver 'simple, secure, and connected' public services by 2030.

Generative AI tools like Microsoft 365 Copilot are already embedded in agency workflows, promising efficiency gains in classification and retrieval.

However, the National Archives of Australia warns that AI-generated records introduce new complexities for retention and disposal.

Agencies must determine how long to keep AI-created content and ensure transparency in automated decision-making processes.

Cybersecurity: The Hidden Compliance Risk

Cyber threats are no longer an IT-only concern - they are an information governance issue. The Australian Signals Directorate's Annual Cyber Threat Report 2024–25 reveals a 23% increase in ransomware incidents targeting government and critical infrastructure.

Data breaches often involve unstructured information repositories, making poor records management a direct contributor to risk exposure.

Adding to the challenge, a 2025 report by Australian Cyber Security Magazine found that 78% of organisations experience cybersecurity burnout, driven by increased threat activity and complex compliance demands.

Alarmingly, 32% of organisations admit to 'shadow AI' use - employees deploying unapproved AI tools that access sensitive data, creating governance blind spots

The New Zealand Perspective

New Zealand agencies face similar pressures. PwC's Global Compliance Survey 2025 notes that many Kiwi organisations still adopt a reactive approach to compliance, addressing issues only after breaches occur.

The report urges a shift toward proactive compliance, embedding governance into strategic planning and leveraging technology for real-time monitoring.

Three Strategic Imperatives for Records Managers

- Embed AI Governance into Records Management: Develop policies for AI-generated records, including retention schedules and audit trails. Ensure transparency and human oversight in automated processes.
- Integrate Cybersecurity with Information Governance: Treat cybersecurity as a compliance issue. Implement data minimisation strategies and enforce systematic disposal to reduce breach impact.
- Adopt Proactive Compliance Frameworks: Move beyond reactive audits. Use maturity assessments (like NSW's RMAT) and invest in tools that automate classification, retention, and disposal.

Conclusion: A Call to Action

The question facing records managers - how to balance compliance, security, and operational efficiency in an AI-driven, cloud-first world - is not hypothetical. It is urgent and real.

Agencies that fail to act risk falling behind regulatory expectations and exposing themselves to cyber threats.

The solution lies in integrated governance: aligning records management, cybersecurity, and AI ethics under a unified framework.

This approach not only ensures compliance but builds resilience and trust in an era of rapid technological change.

Brandon Voight is Fellow at Future Government Institute, and Director of Public Sector - OpenText Australia & New Zealand.

Does Your Business Need ISO 27001 Certification?

By Scytale

Cyberthreats have been making headline after headline. We know it, and you know it... Data security is now more important than ever. And when it comes to managing and protecting information, ISO 27001 certification stands as the gold standard for Information Security Management Systems (ISMS). It helps companies secure data, reduce risk, and build credibility with customers and partners.

But here's the question every business should ask: Do we *actually* need ISO 27001 certification?

Well let's break it down.

First, What is ISO 27001 Certification?

ISO 27001 is an internationally recognized standard that guides companies on how to set up, maintain, and continually improve an Information Security Management System (ISMS).

Getting ISO 27001 certified means your company has proven processes and controls in place to keep data protected, from internal leaks to external attacks. Ultimately, it's about showing your stakeholders that data protection is built into the way you operate.

Why Getting ISO 27001 Certified Matters

If you handle customer data, manage internal systems, or store confidential information, ISO 27001 certification helps you stay secure, compliant, and competitive. And here's how:

- **Protect what matters most** – Keep sensitive information safe from breaches, leaks, and unauthorized access.
- **Stay compliant** – ISO 27001 aligns with many other data protection regulations, helping you meet legal and industry requirements.
- **Stand out from the crowd** – Certification is proof to customers and partners that you take security seriously.
- **Reduce risk, systematically** – The standard gives you a structured way to identify, assess, and manage security threats.
- **Be ready for anything** – From ransomware to insider threats, ISO 27001 builds resilience and strengthens your incident response.
- **Grow with confidence** – A globally recognized certification that opens doors to new markets, customers, and partnerships.

So, Who Needs ISO 27001 Certification?

Short answer: any business that handles sensitive data or operates in a regulated environment. Long answer: some industries benefit even more:

- **Tech and SaaS companies** – If your platform processes or stores customer data, ISO 27001 certification builds trust and helps you pass vendor security assessments faster.
- **Financial institutions and FinTechs** – Protecting financial data and meeting regulator expectations is non-negotiable.
- **Government and public sector** – Agencies managing citizen or national data rely on ISO 27001 to meet strict information security standards.
- **Any business with contractual obligations** – If your clients demand proof of information security compliance, ISO 27001 certification checks that box, and then some.

What You Gain From ISO 27001 Certification

- **A clear, consistent framework** – ISO 27001 gives your team a roadmap to manage data security risks efficiently.
- **More trust, fewer questions** – Certification tells customers, investors, and regulators that you've got data security under control.
- **Fewer breaches, lower costs** – Preventing a single data breach can save millions in fines and lost revenue.
- **Ongoing improvement** – The standard pushes you to regularly review and strengthen your controls as threats evolve.
- **Easier compliance mapping** – ISO 27001 integrates seamlessly with other frameworks like SOC 2 and GDPR, making multi-framework compliance simpler.

How to Get ISO 27001 Certified

In a nutshell, getting ISO 27001 certified is a step-by-step process that starts with building an Information Security Management System (ISMS) tailored to the ISO 27001 standard.

That means assessing risks, defining security controls, and putting policies in place. Once your ISMS is ready, you'll run an internal audit to make sure everything checks out.

Then, you'll bring in an accredited certification body for an official audit. If you pass, you'll earn your certification. But you'll need to keep it up with regular reviews.

It might sound like a lot, but with the right support, the whole process becomes a lot smoother and easier to manage.

So, Does Your Business Need ISO 27001 Certification?

In a world where data breaches make headlines daily, ISO 27001 offers something invaluable: confidence. Confidence that your data is protected, your systems are secure, and your business is prepared for the future.

It's the ultimate investment in your company's data security and long-term resilience.

Whether you're a growing startup or an established enterprise, ISO 27001 certification strengthens your security practices, enhances risk management, and builds trust with customers and partners.

The time is now. Ask yourself: Does ISO 27001 certification align with the needs of my company? Chances are high your answer will be "yes."

And the benefits go far beyond just compliance. They strengthen your security, build trust, and give your business a competitive edge.

Originally published here.

WA Schools delete Student Archives

At least 218 WA government schools deleted student records from the School Information System without proper authorisation between 2010 and 2021. The deletions violated the State Records Act 2000 and breached a 2018 disposal freeze on all records relating to childcare and education.

The Western Australia State Records Commission has issued its first special report to Parliament, detailing the unauthorised destruction of student records at the Department of Education.

The Commission found approximately 9,934 confirmed missing student records, with another 41,000 "orphan" records that could not be matched to student names. Recovery attempts from historic backups achieved less than seven per cent success.

"The premature deletion of records and loss of State archives stored in the SIS is a contravention of the Act," the report states.

The deleted records included State archives such as enrolment and admission records, attendance data, addresses, behavioural records and school reports. Under approved retention schedules, these must be kept for 25 years after the student's birthdate.

System Configuration Failures

The Commission determined the losses resulted from poor practices and improper system configuration rather than malicious intent.

"According to the Department, the system was not configured properly, allowing schools to have the ability to delete records in SIS without the relevant approvals," the report states.

Schools could delete records without approval from the Department's Manager of Corporate Information Services. The system also allowed schools to change default reference number settings, causing certain student sequence numbers to be skipped entirely.

The deletions violated the Department's own Records Management Policy, which requires two-step verification for disposal authorisation involving both the principal and Corporate Information Services.

Royal Commission Disposal Freeze Ignored

The deleted records were subject to a disposal freeze issued by the State Records Office on 5 April 2018, following the Royal Commission into Institutional Responses to Child Sexual Abuse.

The freeze requires indefinite retention of "all records related to the care, supervision, education and treatment of children where government employees, contractors, volunteers and outsourced service providers are in contact with children."

The Royal Commission recommended institutions retain child sexual abuse-related records for at least 45 years to allow for delayed disclosure by victims. "These losses may also impact the investigation of historical abuse allegations and other matters," the report notes.

Remediation Actions

The Department discovered the issue in 2021 and immediately locked down schools' ability to delete student records or use the "Purge Leavers" function.

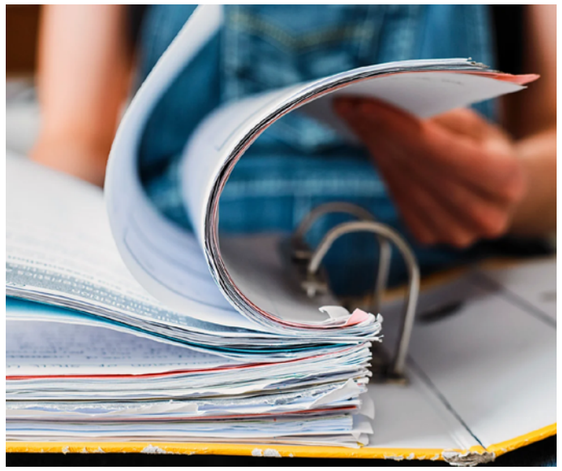

Additional measures included locking student reference numbering settings to ensure sequential numbers, requiring schools to apply to Department ICT staff for deletions with justification, and modifying SIS training to strengthen records management focus.

The Department investigated 200 schools initially, contacting 53 for historic backups. Only nine schools reported backups existed, and just six contained information not already on school servers.

The Department has confirmed the issue did not affect records in the Reporting to Parents system, which retained historical achievement data transferred from SIS since 2010.

The Commission issued two key recommendations for all government organisations subject to the Act.

First, organisations must review business system configurations to restrict record deletion to system administrators and records management staff, not operational users. Systems should monitor all access, modification and deletion of records.

Second, organisations must improve staff awareness of record keeping obligations through onboarding training, regular refreshers, and integration into business system training.

State Records Commission Chair Caroline Spencer, who is also Auditor General for Western Australia, acknowledged the Department self-reported the matter and undertook its own investigation.

"The Commission encourages this practice in the interests of integrity and continuous improvement," Spencer stated in the report's overview.

The Department accepted all recommendations and confirmed it has allocated substantial resources to investigate and remediate the issues. Lessons learned are being embedded in a new student information system designed with enhanced metadata capture, comprehensive audit logging and robust user permissions.

The Department's response noted that over 400 schools have participated in the School Archives Service established in 2019 to support long-term retention of school records following the Royal Commission recommendations.

The full report is available here.

Framing the Digital Health Investment and the Persistent Paper Trap

By Berne Gibbons MAICD

Australia has made significant strides in digital health investment, championing electronic medical records (EMRs) and interoperability initiatives to create a more connected and efficient healthcare system. Yet, beneath the surface, an inconvenient truth persists: vast amounts of critical clinical information remain locked in paper records and siloed EMRs, stifling the promise of digital transformation.

As policymakers, clinicians, and health IT leaders seek to maximize the impact of every dollar spent, it's essential to confront the paradox, our digital ambitions are shackled by the enduring paper trap.

In hospitals across Australia, frontline staff still rely on paper charts, handwritten notes, and printed test results to deliver care. These physical records are often stored in filing cabinets, transported between departments, and, all too frequently, misplaced or delayed. The result is fragmented patient histories, incomplete data at the point of care, and a reliance on memory or guesswork when timely decisions are needed. For patients, this means repeated storytelling, unnecessary delays, and the real risk of errors due to missing information.

The persistence of paper is more than a nuisance, it represents a systemic barrier that undermines clinical workflows, research, and the very interoperability that digital health investments aim to achieve. When vital information is trapped in analog formats, it cannot be efficiently shared, analyzed, or used to drive better outcomes.

Interoperability and Its Limits: The EMR Myth

Electronic Medical Records were heralded as the answer to fragmented healthcare data. While EMRs offer structured digital repositories, their benefits are often overstated. Most EMR systems operate as isolated silos, lacking the seamless connectivity needed for true interoperability. Data sharing between systems is hampered by proprietary formats, inconsistent standards, and limited integration capabilities.

The myth of automatic interoperability has led to complacency. Too often, investments focus on connecting EMRs to each other, ignoring the mountain of paper records and other non-digital sources that remain outside these networks. This oversight perpetuates information gaps and leaves clinicians blind to crucial aspects of a patient's history.

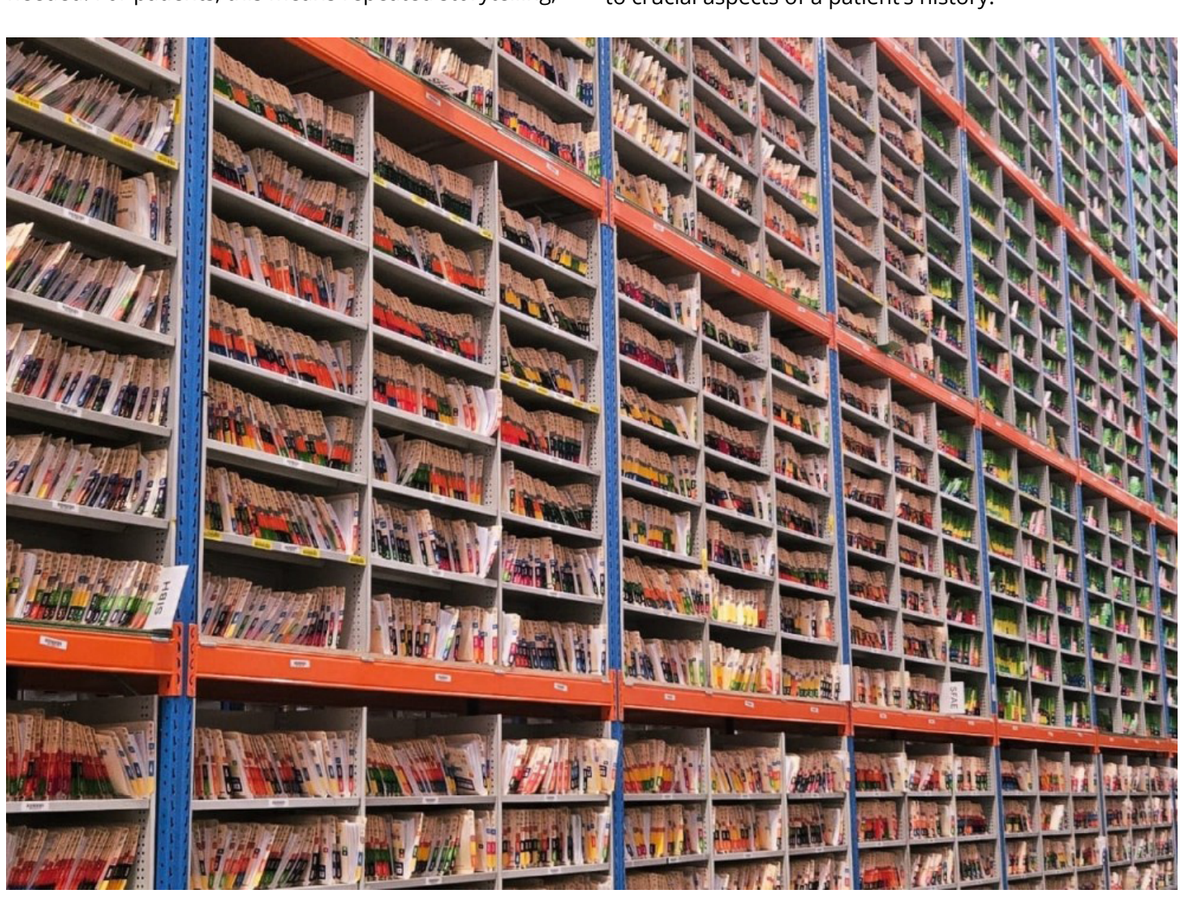

A recent photo taken at a NSW document storagefacility a couple of weeks ago. The workforce that needs to code these (Clinical Coders) need them to be digitised to access remotely.

To break free from the paper trap, digitisation must be front and centre in Australia's health interoperability strategy. Technologies such as Optical Character Recognition (OCR) and Artificial Intelligence (AI) have matured to the point where they can rapidly convert handwritten notes, printed lab results, and legacy documents into structured, searchable digital data. When combined with standardized data formats and robust integration frameworks, these tools unlock previously inaccessible clinical insights.

Globally, leading health systems are reaping the rewards of comprehensive digitisation. The United Kingdom's NHS, for example, has made significant progress in digitizing historical records, using AI to extract and harmonize information for population health studies and real-time clinical support. In the United States, health networks are leveraging OCR and AI to ingest external documents and enable cross-provider data sharing. Even within Australia, pilot programs have demonstrated the feasibility and impact of digitizing pathology reports and discharge summaries - reducing duplication, improving safety, and supporting analytics.

Consequences of Inaction: Duplicated Tests, Hidden Risks, and Wasted Investment

Failing to prioritize digitisation has tangible consequences for patient care. When clinicians cannot see the full picture, tests are repeated, diagnoses delayed, and treatments compromised. Patients face the frustration of retelling their stories and undergoing unnecessary procedures. From a system perspective, millions of dollars are wasted duplicating efforts, while the value of digital investments remains unrealized.

Moreover, hidden risks multiply - adverse drug interactions, missed allergies, and undetected chronic conditions lurk in unseen paper records. Data analytics, population health management, and research are all limited when key information is missing or inaccessible.

A Government-Facing Call to Action: Digitisation as Core Infrastructure

It is time for Australia's health leaders and policymakers to make digitisation a core component of all interoperability initiatives. Funding for digital health must explicitly include the conversion of paper records and legacy documents, not as an afterthought but as essential infrastructure. This means investing in AI-powered OCR, data standardization, and integration platforms that bridge the gap between analogue and digital healthcare.

By embedding digitisation into interoperability planning, the government can ensure that all patient information - regardless of origin - is available, actionable, and secure. Such an approach will maximize the return on digital health investments, empower clinicians with complete data, and drive measurable improvements in patient outcomes.

Conclusion: Bridging the Paper Gap to Realize the Promise of Interoperability

Australia stands at a crossroads. The vision of a truly interoperable, patient-centred digital health system cannot be achieved while vital information remains locked in paper records and siloed EMRs. The path forward is clear: digitisation must be prioritized, funded, and integrated into every aspect of health system transformation. Only then will the promise of digital health interoperability be fulfilled, ushering in a new era of safe, efficient, and connected care for all Australians.

Berne Gibbons MAICD is Chief Strategy Officer, InfoMedix, Member Board of Directors, Standards Australia, and Assoc Professor of Industry - Faculty of Health, UTS. Article originally published here.

NSW Digital Health Record Blowout Beckons

NSW's Single Digital Patient Record (SDPR) project, one of the state government's largest technology initiatives, proceeded with a business case that failed to include critical integration costs and lacked evidence for operational expenses, according to the NSW Auditor-General.

The business case, initially developed in 2021, did not capture the estimated cost of integrating the SDPR system with legacy systems that will remain in use across the health network.

"This integration process is crucial for the successful implementation of the SDPR system and early indicators suggest that these integration costs will be significant," the Auditor-General's report states.

The report warns that "unsupported or unapproved cost estimates increases the risk of budget overruns and misinformed decisions."

The Single Digital Patient Record Implementation Authority (SDPRIA) was established in May 2024 as a division of the Health Administration Corporation to oversee delivery and implementation of the project, which aims to provide clinicians with realtime access to patient medical information from a single source.

The operational cost estimates included provisions for implementation-related expenses across local health districts and in-scope health entities, but "this estimate was not supported by robust documentation due to limited cost information available at the time," according to the report.

The Auditor-General noted that "project costs may be understated by not assessing and identifying all implementation costs during project planning."

The SDPRIA recently awarded a contract worth $A83 million to RLDatix Galen Australia for a statewide data archive solution (see over page),

A performance audit is planned for 2026-27 to assess the efficiency and effectiveness of the project implementation.

$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.

Healthcare provider slashes document processing costs with GenAI

A genetic testing company has demonstrated how generative AI can deliver substantial cost reductions and efficiency gains in document processing workflows, offering insights for information managers grappling with high-volume document challenges.

Myriad Genetics, a Salt Lake City-based provider of genetic testing and precision medicine solutions, partnered with AWS to transform its healthcare document processing pipeline using Amazon Bedrock foundation models. The implementation reduced document classification costs by 77 percent while improving accuracy from 94 to 98 percent.

The company processes thousands of healthcare documents daily across its Women's Health, Oncology and Mental Health divisions. Documents must be classified into categories including test request forms, lab results, clinical notes and insurance records to automate prior authorisation workflows.

Myriad's existing solution combined Amazon Textract for optical character recognition with Amazon Comprehend for document classification. The system routed classified documents to appropriate external vendors for processing based on their identified document class.

Despite achieving 94 percent classification accuracy, the system created significant operational challenges. Processing costs reached 3 cents per page, resulting in monthly expenses of $US15,000 per business unit. Classification latency averaged 8.5 minutes per document, creating bottlenecks in downstream workflows and delaying prior authorisation submissions.

The existing system's limitations extended beyond classification. Key information extraction remained entirely manual due to the complexity of medical documents. Staff needed contextual understanding to differentiate critical clinical distinctions such as 'is metastatic' versus 'is not metastatic' and to locate information including insurance numbers and patient details across varying document formats.

This manual processing burden was substantial. The Women's Health business unit alone required up to 10 full-time employees contributing 78 hours daily to extract key information from documents. The workload created scalability constraints and operational bottlenecks as document volumes increased.

Myriad needed a solution to reduce document classification costs while maintaining or improving accuracy, accelerate document processing to eliminate workflow bottlenecks, automate information extraction for medical documents, and scale across multiple business units and document types.

The new solution uses AWS's open-source GenAI Intelligent Document Processing Accelerator with Amazon Nova foundation models. Amazon Nova Pro handles document classification, while Amazon Nova Premier manages complex information extraction requiring advanced reasoning capabilities.

Processing speed improved by 80 percent, reducing classification time from 8.5 minutes to 1.5 minutes per document. The automated key information extraction achieved 90 percent accuracy, matching human evaluator baseline performance, while processing documents in approximately 1.3 minutes each.

Implementation challenges centred on handling complex medical documents with visual ambiguity. Checkbox fields required distinguishing between different marking styles including checkmarks, crosses and handwritten marks. Documents contained overlapping marks and content spanning multiple fields.

AWS Data Scientist Priyashree Roy and her team addressed these challenges through several optimization techniques. They enabled Amazon Textract's specialized tables and forms features to improve optical character recognition discrimination between selected and unselected checkbox elements.

The team implemented a multimodal approach that sent both document images and extracted text to the foundation model, enabling simultaneous analysis of visual layout and textual content. Few-shot learning provided example document images paired with expected extraction outputs to guide the model's understanding of various form layouts.

For particularly complex extraction scenarios, the team used Amazon Nova Premier with chain of thought reasoning, having the model work through extraction decisions step-by-step before making final determinations.

Prompt engineering proved critical to achieving high accuracy. The team used document samples from each class with Anthropic Claude Sonnet 3.7 on Amazon Bedrock with model reasoning enabled. The model identified distinguishing features between similar document classes, which Myriad's subject matter experts refined.

Classification Strategies

Format-based classification strategies used document structure and formatting as key differentiators. Lab reports contain numerical results organized in tables with reference ranges and units, while test results present findings in paragraph format with clinical interpretations.

Negative prompting techniques resolved confusion between similar documents by explicitly instructing the model what classifications to avoid. Test request forms were frequently misclassified as test results due to confusion between patient medical history and lab measurements. Adding exclusionary language to classification prompts improved classification accuracy by 4 percent.

Myriad plans a phased rollout beginning with document classification in the Women's Health business unit, followed by Oncology and Mental Health divisions. The company projects annual savings of $US132,000 in document classification costs.

The solution reduces each prior authorisation submission time by 2 minutes. Specialists now complete orders in four minutes instead of six minutes due to faster access to tagged documents. This improvement saves 300 hours monthly across 9,000 prior authorisations in Women's Health alone.

Martyna Shallenberg, Senior Director of Software Engineering at Myriad Genetics, said the partnership with AWS GenAI Innovation Centre delivered measurable business impact. She cited improved performance and accuracy alongside projected savings of more than $US10,000 per month.

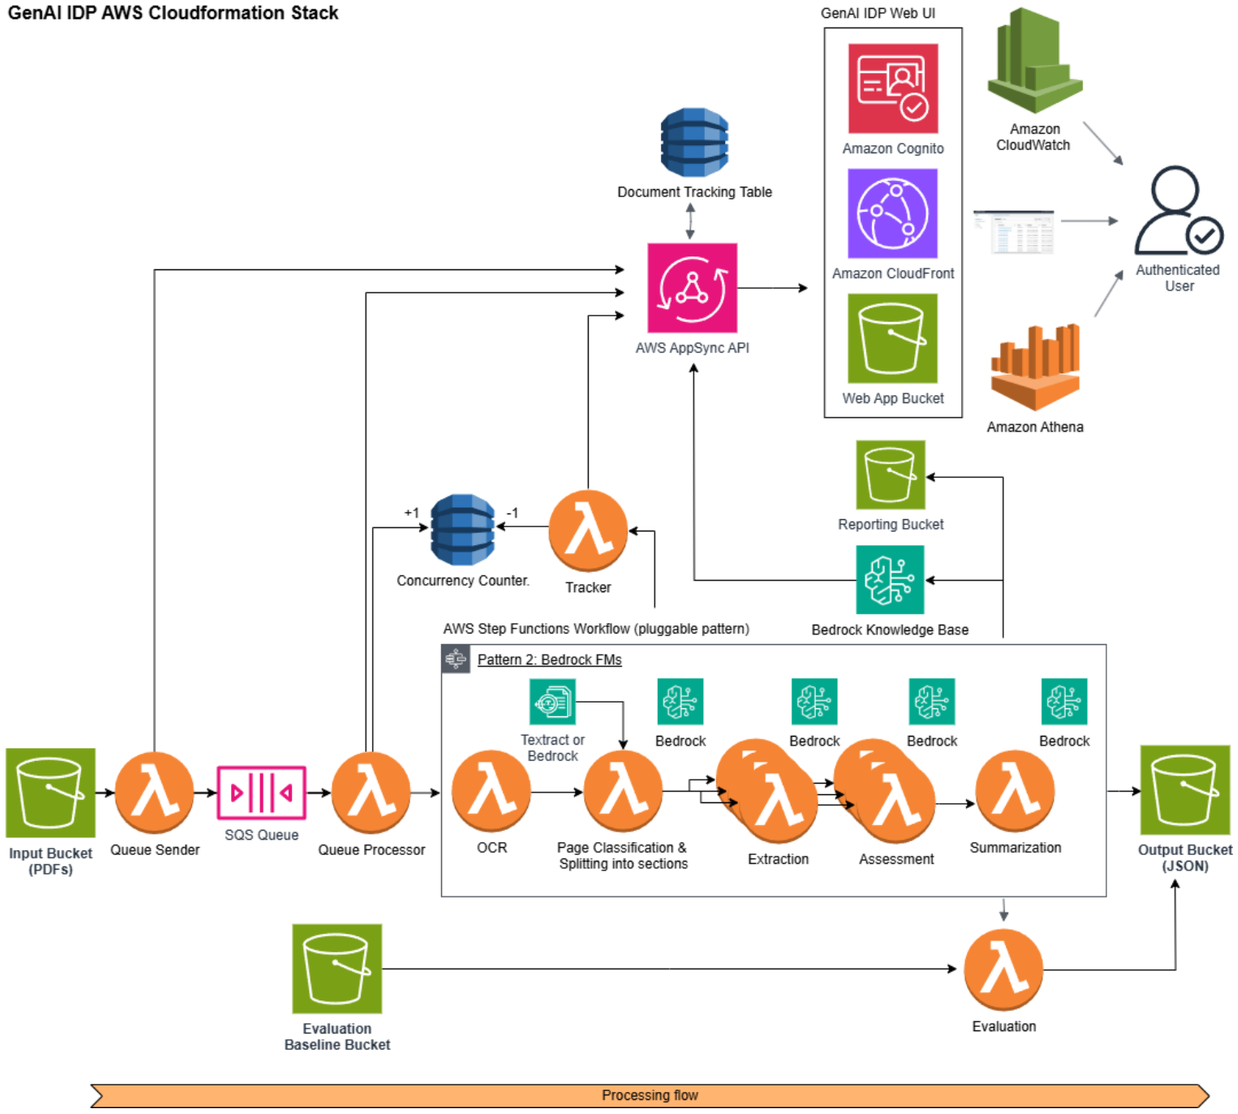

The diagram illustrates the solution architecture, showing the default Bedrock Data Automation workflow (Pattern-1).

The GenAI Intelligent Document Processing Accelerator provides a serverless architecture that converts unstructured documents into structured data. The accelerator processes multiple documents in parallel through configurable concurrency limits. Its built-in evaluation framework lets users provide expected output through the user interface and evaluate generated results to customize configuration.

The accelerator offers one-click deployment with three pre-built patterns optimized for different workloads.

Pattern 1 uses Amazon Bedrock Data Automation with out-of-the-box features and straightforward per-page pricing. Pattern 2 uses Amazon Textract and Amazon Bedrock with Amazon Nova or Anthropic Claude models, ideal for complex documents requiring custom logic. Pattern 3 uses Amazon Textract, Amazon SageMaker with fine-tuned models for classification, and Amazon Bedrock for extraction.

Myriad selected Pattern 2 to meet requirements for low cost while offering flexibility to optimize accuracy through prompt engineering and model selection. The pattern offers no-code configuration, allowing customization of document types, extraction fields and processing logic through configuration editable in the web interface.

The open-source GenAI IDP Accelerator is available for organizations to deploy and test in their environments.

AWS provides detailed documentation on accelerating intelligent document processing with generative AI at https://aws.amazon.com/blogs/machine-learning/accelerate-intelligent-document-processing-with-generative-ai-on-aws/

The GenAI IDP Accelerator is available at http://www.amazon.com/genai-idp-accelerator

This article is based on a case study published by AWS:

The Evolution of Data Governance in the Age of AI

By Chad Barendse

I was recently asked what impact AI will have on data governance and it got me thinking. Most teams I see still run governance the old way: central policy group, registers, manual checks. That worked when data moved slower. But the pace of delivery and the rise of AI has changed everything.

For decades, data governance meant a central team setting policy, running councils, approving access, and manually keeping registers of sensitive data. That approach matched the technology and risk profile of the time: slower release cycles, fewer systems, and compliance paperwork.

Now everything about how we build and use data is changing:

- Data is infrastructure and code.
- AI is everywhere.
- Delivery is rapid.

Manual, centralised governance can't keep up. It must become engineered and embedded, the same shift security made with DevSecOps.

When I last wrote about the changing role of the Data Governance Manager, the question wasn't whether data governance still matters. It was clear it does. The question was **how** it should work in a world defined by fast delivery and AI.

For decades, data governance meant a central team setting policy, running councils, approving access, and manually keeping registers of sensitive data. That approach matched the technology and risk profile of the time: slower release cycles, fewer systems, and compliance paperwork.

Manual, centralised governance can't keep up. It must become **engineered, embedded**, the same shift security made with DevSecOps.

AI Is Changing Both Sides of the Equation

AI isn't just something we have to govern. It's also transforming how data governance gets done.

- **Automating the heavy lifting —** AI scans for personal or sensitive data, tags and classifies, maps lineage, and flags quality issues.
- **Guiding and assisting —** Large language models help developers and analysts understand compliance rules and suggest fixes.
- **Guardian agents —** AI enforces policies in real time: trigger deletion, revoke unsafe access, block risky models, validate controls, and improve data quality.

This is what finally makes "governance as code" practical.

From Policy to Engineered Governance

We've reached a point where the old ways can't keep up. AI changes the risk surface and gives us new tools. The next step isn't to abandon data governance, it's to rebuild it around how modern products and AI are actually delivered.

That means moving from a central team writing and checking policy **after the fact** to an operating model where controls are engineered into the flow of work, assisted by AI, and owned by the people closest to the data and models.'

The New Operating Model

Today, most data governance teams look roughly the same: a manager or lead who sets policy, a data quality analyst or two, someone keeping the catalog alive, and some privacy support. It's centralised, manual, and slow — and it breaks once you have dozens of data products and AI models shipping every week.

What's emerging instead is an operating model built for speed and automation:

Executive Accountability stays — the Chief Data & AI Officer still owns data and AI risk appetite and reports to the board.

Delivery moves closer to the work — product owners own specific governance capabilities (quality, protection, discovery) and work directly with squads building data and AI products.

AI does the grunt work — scanning, tagging, risk detection, and enforcing basic rules.

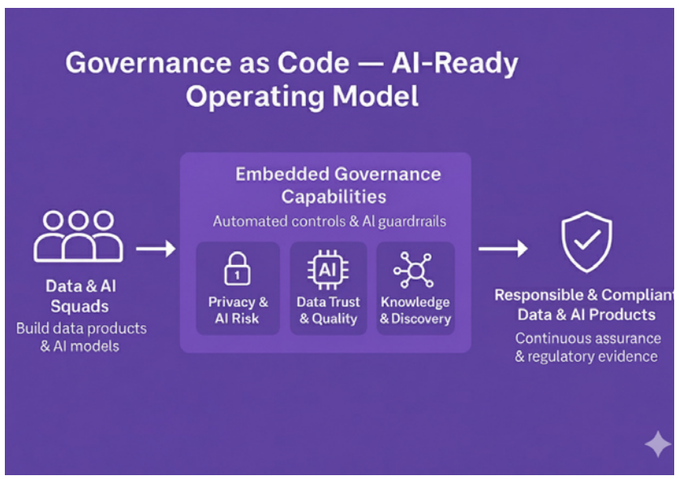

Enablement gets serious — literacy and adoption aren't side projects; they're run as a product with a lead and clear outcomes.

The Takeaway

Data governance isn't dying - it's evolving. The days of manual registers and council approvals are done. Governance now needs to be engineered and automated.

That means fewer policy writers and more **product owners and engineers,** building controls and automation into the delivery flow. AI isn't replacing governance - it's forcing it to become faster, smarter, and closer to the work.

Chad Barendse is Co-founder of DGX Group, a consultancy dedicated solely to Data and AI Governance. Learn more at dgx.group or follow on LinkedIn for practical insights on modern data and AI governance.

Making Data Governance Work With, Not Against, Your Teams

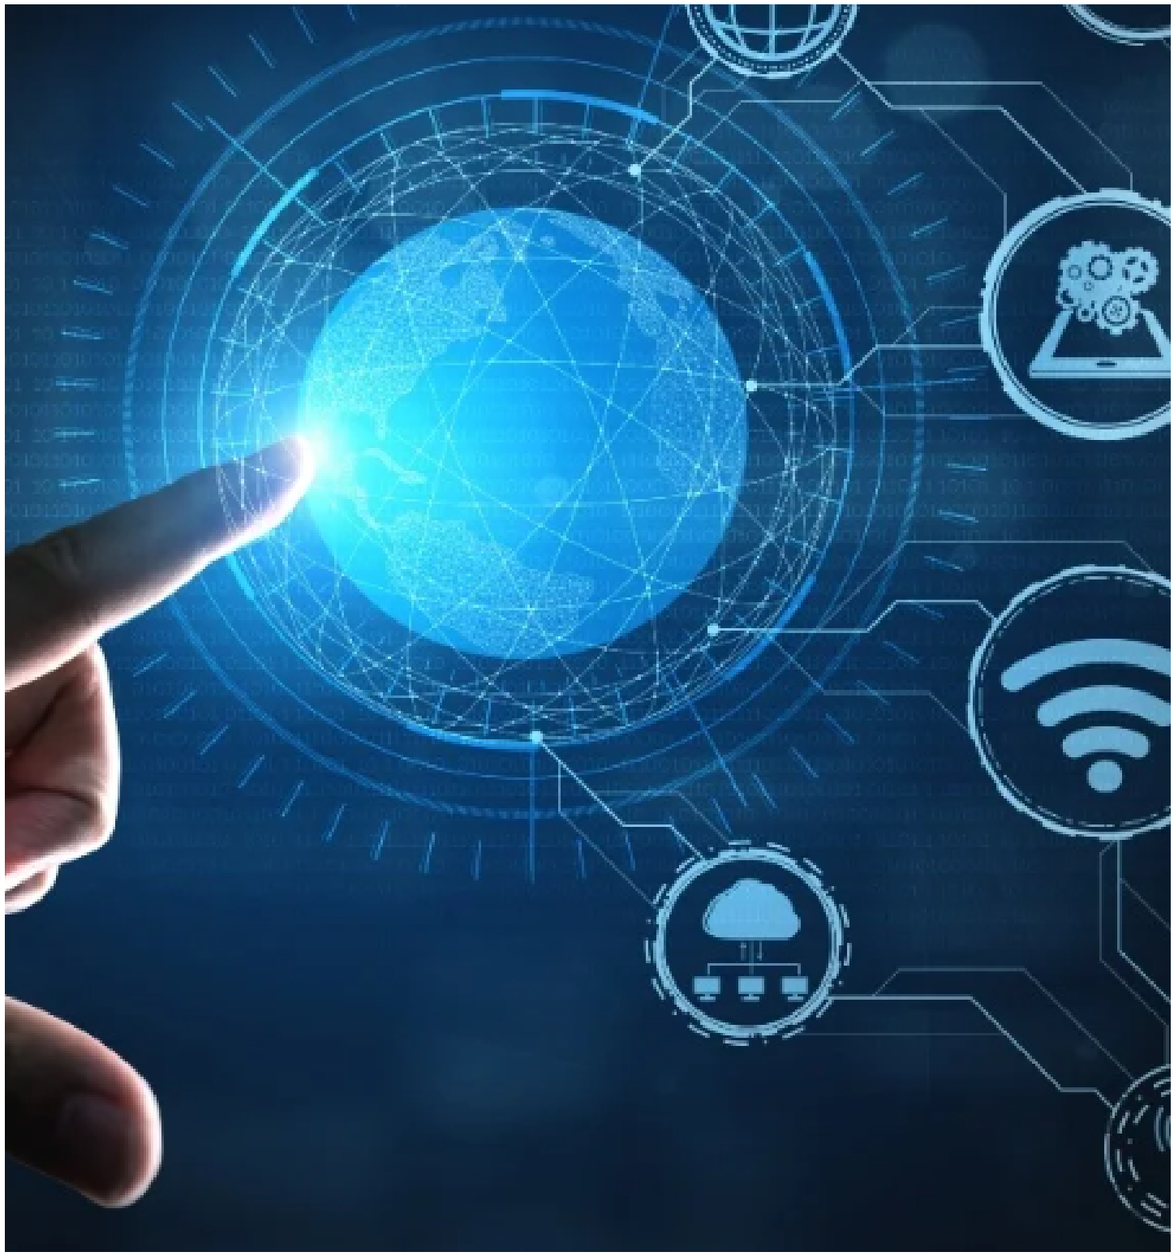

By Nicola Askham

If you've been following data governance discussions lately, you'll have noticed a troubling pattern: despite significant investments in governance teams, data catalogs, and policy frameworks, organisations are still struggling with the same fundamental challenges. Poor data quality, compliance violations, and that persistent lack of trust in enterprise data.

Sound familiar? You're not alone.

The uncomfortable truth is that most data governance approaches are fundamentally broken—not because the concept is wrong, but because we've been implementing it backwards.

What Traditional Data Governance Gets Wrong

Let me start by addressing the elephant in the room: traditional data governance is reactive. We've built entire frameworks around fixing problems *after* they occur, rather than preventing them from happening in the first place.

Here's what this typically looks like in practice:

- Data gets created and pushed to production without proper metadata
- Data stewards scramble to document and classify assets after the fact
- Quality issues are discovered downstream when reports fail or decisions go wrong
- Data governance teams spend their time in endless catch-up mode, always one step behind

This approach creates what I call the "governance gap"—the dangerous space between when data is created and when it's properly governed.

During this gap, data consumers lose trust, compliance risks multiply, and the entire data governance program starts to feel like an expensive afterthought.

The Knowledge Hand-off Problem

One of the biggest issues I see is the problematic knowledge transfer between domain experts, data engineers, and data stewards.

Think about it: the person who best understands the business context of the data (the domain expert) isn't the same person building the data pipeline (the data engineer), who also isn't the same person responsible for cataloguing it (the data steward).

Each hand-off is an opportunity for critical context to get lost. By the time your data reaches production, much of its business meaning has been diluted or completely misunderstood.

Absolutely not ideal, is it?

Introducing Governance Shift Left: A Better Way Forward

Here's where things get interesting. What if instead of treating governance as a separate, downstream activity, we embedded it directly into the data engineering process from the very beginning?

This is the essence of Governance Shift Left - a proactive approach that integrates governance practices into the earliest stages of the data lifecycle, particularly during the software implementation phase when data pipelines are being built.

The concept isn't entirely new (software development has been "shifting left" on testing and security for years), but its application to data governance represents a fundamental paradigm shift.

The Four Pillars Of Governance Shift Left

Governance Shift Left is built on four core principles:

1. Lifecycle Alignment Metadata, code, and data should follow the same development lifecycle. They're all part of the business value you're creating, so why manage them separately?

2. Ownership Your data engineering team becomes directly accountable for compliance, not just data delivery. They adopt governance policies as part of their standard development process.

3. Policy as Code Governance policies are no longer guidelines—they're automatically enforced through code and cannot be bypassed. This transforms abstract policies into concrete, executable rules.

4. Transparent Documentation Policies should be documented, accessible, and self-explanatory. A good policy explains not just what to do, but why it exists and what the trade-offs are.

Why This Approach Works

When you align data documentation with the software development lifecycle, you can apply the same quality gates you use for code before it goes into production. The benefits compound quickly:

Improved Time to Market: No more waiting for separate governance teams to catch up with your data initiatives. Quality and compliance are built in from day one.

Reduced Manual Effort: Your data catalog automatically stays aligned with governance policies, eliminating the need for manual data entry and reducing errors.

Enhanced Trust: When data and metadata are created together and never fall out of sync, data consumers can rely on what they find in your catalog.

Lower Costs: Fewer manual checks, less rework, and reduced maintenance costs as quality issues are prevented rather than fixed.

Making It Practical: Data Contracts and Policy Automation

Two key enablers make Governance Shift Left practical rather than just theoretical:

Data Contracts serve as software-defined agreements that include technical schemas, business metadata, SLAs, and quality expectations. These become artefacts produced by your data teams, enabling governance enforcement at deployment time.

Policy as Code provides the ability to build automated quality gates for metadata and enforce them during your CI/CD process. These can be sophisticated - checking if semantics align with your business glossary or ensuring compliance with industry regulations.

Your Next Steps

If you're ready to move beyond reactive governance, here's what you should do:

- **Start Small**: Identify one critical data pipeline and implement basic data contracts
- **Align Teams**: Bring your data governance and engineering teams together to define policies that can be automated
- **Implement Quality Gates**: Add metadata validation to your CI/CD pipeline
- **Measure Impact**: Track the reduction in downstream quality issues and governance effort

The shift won't happen overnight, and you'll need buy-in from both technical and business stakeholders. But the alternative - continuing with reactive, resource-intensive governance - simply isn't sustainable as data volumes and complexity continue to grow.

The Bottom Line

Traditional data governance assumes that good governance happens *to* data after it's created. Governance Shift Left recognises that good governance happens *with* data as it's being created.

This isn't just about improving your governance program—it's about fundamentally changing how your organisation thinks about data responsibility and quality.

The question isn't whether you can afford to make this shift. It's whether you can afford not to.

Originally posted on www.nicolaaskham.com

Active Archives meet Soaring Storage and Energy Demand

Organisations must shift from static, passive data archives to intelligent, accessible repositories that unlock archival data value, according to a new report released by the Active Archive Alliance.

The report – "*Preparing for Tomorrow's Expanding Storage Challenge with Active Archive*" – reveals how active archives are transforming data management strategies across enterprise sectors managing compliance, artificial intelligence, and digital transformation initiatives.

The shift comes as data centre electricity demand accelerates sharply, driven by artificial intelligence workloads. In the United States, power consumption by data centres is projected to account for almost half of electricity demand growth between 2025 and 2030, the report states.

Electricity demand from AI–optimised data centres is projected to more than quadruple by 2030, reaching 945 terawatt–hours (TWh) – slightly more than Japan's entire electricity consumption today.

For Information Managers and Digital Transformation Managers, the implications are profound. Active archives enable reliable, online, and cost–effective access to data throughout its lifecycle whilst managing energy consumption and compliance obligations. Approximately 80 percent of an organisation's digital data falls into the low–activity or inactive category, making active archives essential for managing this fastest–growing data class.

The Active Archive Alliance describes an active archive as a strategic approach to intelligently manage data throughout its lifecycle, allowing organisations to balance immediate access to archival data with storage efficiency.

Intelligent Data Management Software (IDMS)

Unlike traditional cold archives – often static and difficult to access – active archives integrate two or more storage technologies behind Intelligent Data Management Software (IDMS) to provide seamless management of archive data in a single virtualised storage pool.

"Demand for enterprise storage capacity is accelerating rapidly, driven by massive AI–fuelled data growth," said Rich Gadomski, co–chairperson of the Active Archive Alliance and director of channel sales at FUJIFILM North America Corp.

"Active archives enable smart storage optimisation that offers the right balance of accessibility, performance and energy efficiency by moving data between storage tiers based on access frequency and business value."

Energy efficiency represents a critical advantage for enterprise technology leaders. By migrating low–activity or inactive data from hard disk drives (HDDs) to magnetic tape through an active archive strategy, organisations can reduce carbon emissions by up to 97 percent compared to all–HDD solutions.

A study comparing three scenarios – all–HDD, all–tape, and an active archive strategy moving 60 percent of low–activity HDD data to tape – demonstrated a 58 percent reduction in carbon emissions over a 10–year period.

The report highlights critical challenges confronting organisations implementing long–term archival strategies. Limited metadata for unstructured data complicates efficient data retrieval and AI training processes. Organisations face significant risks of format obsolescence – a document saved in 2025 may not be readable in 2075 due to software ecosystem changes, vendor discontinuation, or lack of backward compatibility.

Technology Managers and CTOs must account for hardware refresh cycles. Hard disk drives typically last 4–6 years before replacement or failure. Data retained for 100 years would require at least 10 migrations to newer media, each introducing risks of data loss, corruption, and increased operational costs. A recent Experian study revealed that 64 percent of data migration projects analysed exceeded budgets, and only 46 percent delivered on time.

Data migration challenges underscore the need for systematic planning. The report identifies four key challenges and risks: making archival data accessible at ingest through metadata classification and indexing; managing long–term archival storage infrastructure; ensuring only potentially needed archive data is stored; and ensuring security and availability of archival data.

LTO technology roadmap

In 2025, the LTO Consortium launched LTO–10, the tenth generation of the industry's most widely adopted tape storage format. The LTO technology roadmap extends through 14 generations, with LTO–14 delivering up to 1,440 TB (1.44 petabyte) compressed per cartridge.

A record 176.5 exabytes of total tape capacity (compressed) shipped in 2024, representing 15.4 percent growth over 2023. This growth reflects evolving infrastructure requirements driven by AI, machine learning, and unstructured data growth – alongside shifts toward lower–cost hybrid cloud environments.

For hybrid cloud strategies, active archives bridge on–premises and cloud environments, enabling seamless long–term storage and access to archival data. Frequently accessed data can remain on–premises or in high–performance cloud tiers, whilst less frequently accessed data automatically moves to low–cost cloud storage. A unified global namespace provides a single logical view of data across both on–premises and cloud locations.

The report introduces a transformative secondary storage model comprising three distinct tiers. The Active Archive tier (Write–Once, Read–Many) provides online access to dynamic archival data. The Traditional Archive tier (Write–Once, Read–Seldom if Ever) supports lower–activity, big data and cold archives. The emerging Deep Archive tier (Write–Once, Read–Never) targets permanent, rarely accessed, dark data – often serving as a golden, immutable master copy.

The full report is available here.

Three Top Cybersecurity Projects to Prioritize: Gartner

By Fadeen Davis, Garner Inc.

If the expression "May you live in interesting times" describes your role, you may be a security and risk management (SRM) leader. Between resource constraints caused by rapid-fire tariff announcements, pressure to expand your purview and a perpetual need to catch up with the latest cybersecurity threat, your best bet is to prioritize initiatives that align with the organization's broader objectives.

Staffing shortages and tight budgets are putting comprehensive security solutions out of reach for many SRM leaders. We spotlight three 12-month projects that optimize resources and engage stakeholders.

The SRM leader's role is changing as businesses are taking on more risks, responding to the macro-economic conditions and the evolving threat landscape. This change has led to additional responsibilities for the SRM leader such as cyber-physical systems security (e.g., operational technology, the Internet of Things, enhancing data governance frameworks, ensuring robust business continuity strategies and facilitating the adoption of AI technologies).

SRM leaders must also navigate new data and AI compliance regulations, a constant barrage of new threats and increasing complexity in IT. Uncertainty about tariffs tops it all off.

SRM leaders can make measurable progress in the next 12 months by:

- Introducing the safe use of GenAI by embedding cybersecurity considerations into the GenAI governance framework
- Adopting risk-reducing practices and technologies to unstructured data management
- Implementing advanced storage governance and AI-driven threat detection

Embed cybersecurity considerations into GenAI governance

As organizations integrate and commoditize GenAI in operations, SRM leaders must make it their business to establish strong security guardrails that protect the technology and its use. GenAI is deployed via software embeds and emerging structures (e.g., agentic AI), so policies must be updated routinely to address emerging capabilities and security implications.

Approach this initiative with the assumption that GenAI capabilities are embedded in most software and that users cannot detect the presence of GenAI functions in applications. This mindset covers the interests of the organization while seeking to preserve stakeholder trust.

SRM leaders should create short-, medium- and long-term goals to define and implement GenAI governance:

- **Short term.** Augment governance artifacts with GenAI guidance by creatingGenAI policy, standards, procedures and usage guidelines to cover GenAI-specific security issues.
- **Medium term.** Monitor for compliance and identify governance deficiencies.
- **Long term.** Collaborate with business leaders to address gaps produced by evolving GenAI adoption and deployment models.

> This year, prioritize cybersecurity projects that reduce organizational risk and increase resilience and can be completed in 12 months.

Prepare unstructured data for GenAI adoption

The top reason for failed GenAI deployments is a lack of GenAI-ready data. Organizations that use retrieval-augmented generation pipelines for their GenAI applications need access to unstructured data, which comprises 70% to 90% of the data in today's organizations.

As vendors race to address demand for augmented data management, SRM leaders must establish data governance programs that enforce security measures to prevent GenAI from allowing unauthorized access to sensitive data. This means discovering, classifying and cataloguing unstructured data, optimizing access entitlements and protecting sensitive content by redacting or concealing it.

SRM leaders should:

- Implement a data governance program to address the specific challenges and benefits of integrating GenAI into the organization.
- Reduce unnecessary data and streamline data preparation for AI applications.
- Make sure data discovery tools, data classification and handling instructions are ready for AI use.
- Ensure compliance with data residency requirements, especially GenAI services that involve cross-border data transfers.
- Maintain clear visibility and transparency in the use of AI in standardized business applications and custom AI applications that require large language model (LLM) training data.

Enhance the cybersecurity of data with cyberstorage

Traditional infrastructure security layers that focus on network and endpoint are no longer enough to prevent cyberattacks and data loss. By 2029, 100% of storage products will include cyberstorage capabilities focused on active defence beyond recovery from cyber events, up from 20% in early 2024. Cyberstorage adds a security layer at the storage level – a proactive bulwark against advanced cyberthreats.

Storage teams may be reluctant to step outside their comfort zone. However, as cyberattacks increase and gain visibility at C-level, storage operators must accept responsibility for implementing cyberstorage capabilities.

Ultimately, if the security layers at the network perimeter or application levels fail, the storage administrator is responsible for identifying what data to recover and from what point in time, while ensuring access to critical data.

Enterprise storage solutions vendors have started to offer a range of security features that align with the U.S. National Institute of Standards and Technology (NIST) Cybersecurity Framework's six pillars: govern, identify, protect, detect, respond and recover.

What is cyberstorage? Cyberstorage actively defends storage systems and data against cyberattacks through prevention, early detection and attack blocking. It also supports data analytics, forensic attack analyses and storage-specific recovery capabilities. Cyberstorage can be an enterprise storage product with comprehensive features, a platform-native service offering with integrated solutions or a collection of standalone products that augment storage vendors with cyber protection capabilities.

How are tariffs impacting cybersecurity? Organizations are facing uncertainties due to the speed at which tariffs are being announced and applied. These have the potential to impact security planning and operations in different ways, such as delaying deployment of critical security infrastructure

Fadeen Davis is a Senior Principal Advisor with Gartner's Security Leaders Practice. Be part of the biggest gathering of security and risk leaders at the Gartner Security & Risk Management Summit 2006 on March 16-17, Sydney.

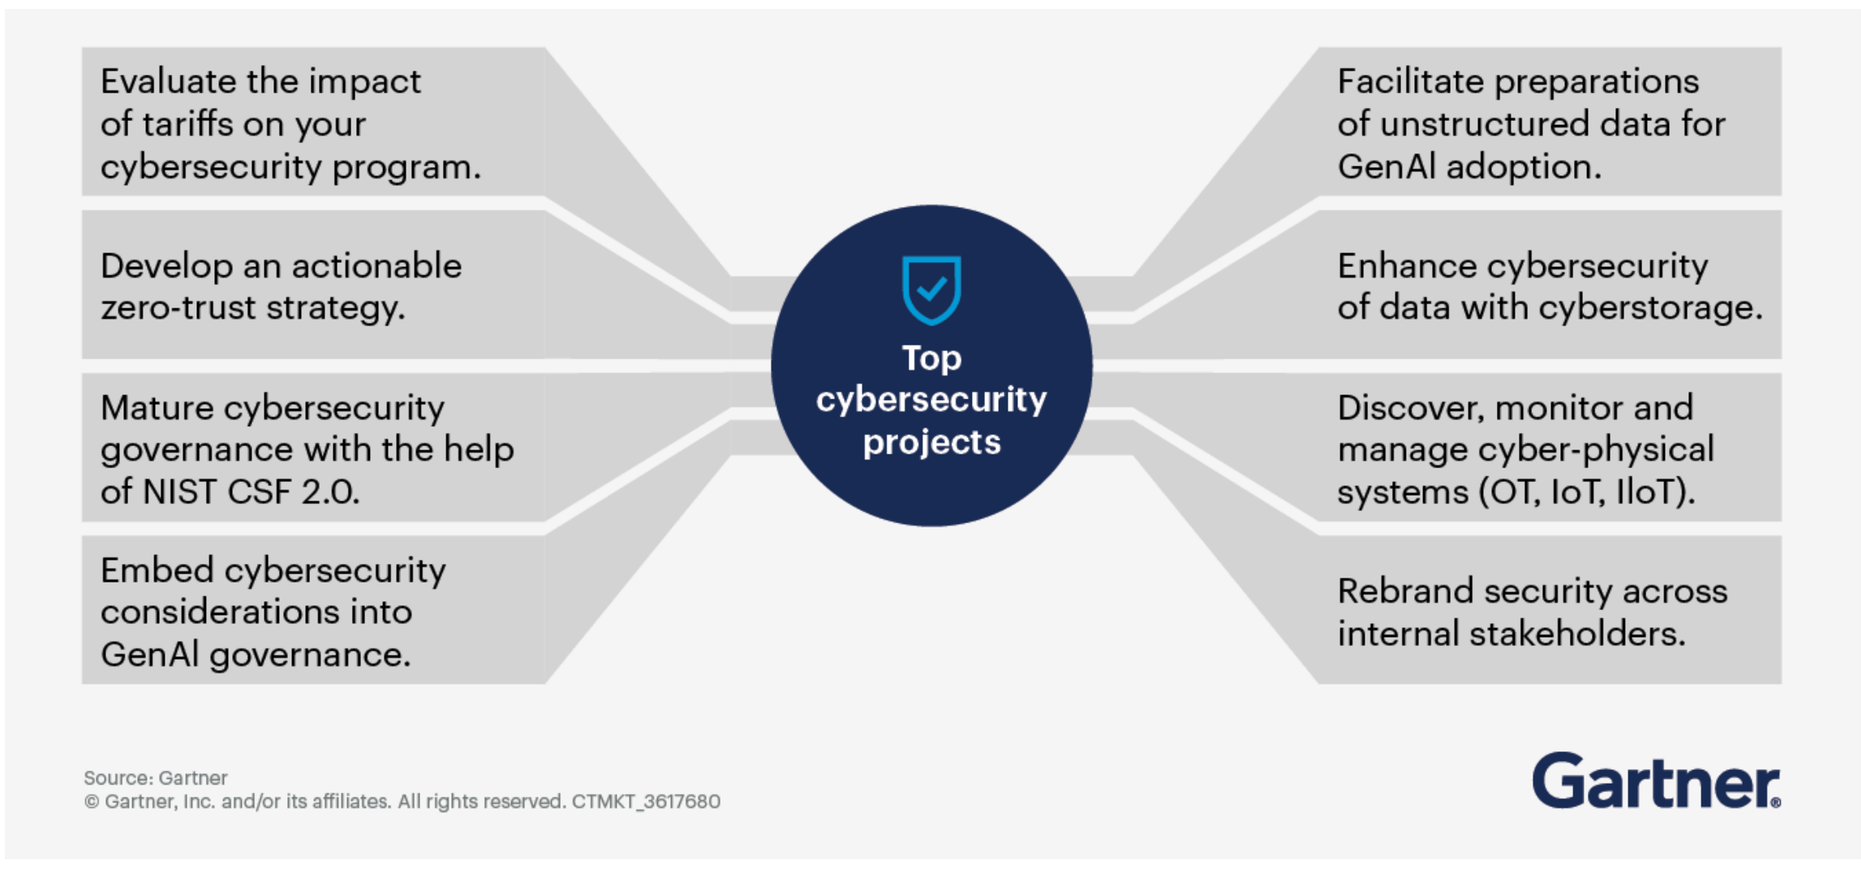

EzeScan is one of Australia's most popular production capture applications and software of choice for many Records and Information Managers. This award winning technology has been developed by Outback Imaging, an Australian Research and Development company operating since 2002. Solutions range from centralised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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

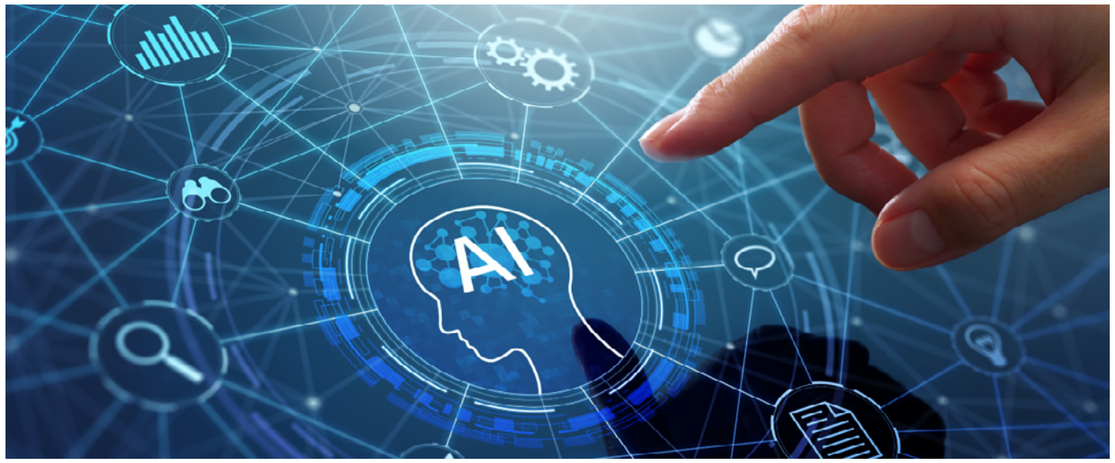

Registry Solution Tackles Agentic AI

Designed to provide organisations with centralised visibility and control over autonomous AI tools, AvePoint has launched AgentPulse Command Centre, an AI agent registry.

The platform addresses growing security and cost management challenges as enterprises deploy AI agents across their operations. AgentPulse enables IT and security teams to track which agents are active, identify users, and monitor access to sensitive data.

The launch follows research indicating significant AI-related security risks. According to data cited by AvePoint, 75 per cent of organisations using AI experienced data breaches in the past year. Gartner predicts 40 per cent of agentic AI projects could be abandoned by 2027 due to inadequate risk controls.

Agentic AI systems operate with greater autonomy than traditional chatbots, making independent decisions and taking actions within business systems. This independence creates new security challenges beyond those addressed by legacy IT controls.

Prompt injection attacks represent a particular concern, where malicious actors manipulate AI decision-making processes by embedding harmful instructions in external data sources. Recent research has identified these attacks as systemic challenges affecting AI-powered browser agents across the industry.

AgentPulse shows which AI agents generate the most activity and access large data volumes. Organisations can use this information to identify unnecessary agents driving unexpected costs and adjust or remove them.

The tool also identifies oversharing instances where agents have excessive access to sensitive files. Teams can remediate these issues directly within the Confidence Platform.

“As organisations rapidly deploy AI agents to drive productivity, they’re discovering that these powerful tools introduce significant security and data governance challenges that legacy IT controls weren’t designed to address,” said Jeremy Thake, Chief Architect, AvePoint.

“As part of the Confidence Platform, AgentPulse delivers comprehensive visibility so IT and security teams can confidently scale agentic AI and further innovation, without compromising data security or racking up extra costs.”

John Peluso, Chief Technology Officer at AvePoint, noted the financial implications extend beyond licensing costs. Without oversight, organisations face unexpected charges from high-activity agents, storage costs from redundant data, and potential breach expenses.

The solution follows an October update to the Confidence Platform that streamlined compliance enforcement for AI agents. AgentPulse supports multiple agent types including Microsoft 365 Agents Toolkit, Microsoft 365 Copilot Studio and Copilot Studio Lite, SharePoint, Azure AI Foundry, ISV Store, and more.

Gartner analysts have warned about “agent washing,” where vendors rebrand existing tools as agentic AI without delivering autonomous capabilities. The analyst firm estimates that only about 130 vendors out of thousands claiming agentic solutions actually offer genuine agentic features.

AvedPoint AgentPulse is available in private preview. Organisations can request access through AvePoint’s website.

https://www.avepoint.com/solutions/ai-agent-governance-and-security

Joint push for enterprise AI agents

Snowflake and Anthropic have signed a $US200 million agreement to embed AI agents into enterprise data platforms. The multi-year deal integrates Anthropic’s Claude models directly into Snowflake’s governed environment.

This partnership targets CIOs struggling to move generative AI from experimental pilots to secure production workflows.

This partnership competes directly with Salesforce’s “Agentforce” and Microsoft’s Copilot Studio, which also promise secure, data-grounded AI agents. The unique selling proposition here is data gravity: bringing the model to where the data already lives (Snowflake), rather than moving data to the model.

The agreement establishes a joint go-to-market initiative focused on deploying AI agents within large global enterprises.

Snowflake Intelligence now uses Claude Sonnet 4.5 to analyse structured and unstructured data via natural language.

The integration aims to solve data leakage concerns by keeping AI reasoning within Snowflake’s security perimeter.

Snowflake Horizon Catalog provides the governance controls necessary for regulated industries like finance and healthcare.

Enterprises can deploy agents that understand customer data context while maintaining strict compliance standards.

“Enterprises have spent years building secure, trusted data environments, and now they want AI that can work within those environments without compromise,” said Dario Amodei, CEO and Co-Founder of Anthropic. “This partnership brings Claude directly into Snowflake, where that data already lives. It’s a meaningful step toward making frontier AI genuinely useful for businesses.”

The partnership introduces capabilities for multimodal analysis across text, images, audio, and tabular data.

Customers can access these features using SQL through Snowflake Cortex AI Functions.

Snowflake Cortex Agents allow organisations to build custom, production-ready data agents.

These agents are designed to retrieve information and reason over data with verifiable accuracy.

Snowflake claims Claude is the only frontier model available across Amazon Bedrock, Google Cloud Vertex AI, and Microsoft Azure.

The vendor states Claude achieves over 90 per cent accuracy on complex text-to-SQL tasks.

This figure relies on Snowflake’s internal benchmarks and has not been independently verified.

Snowflake CEO Sridhar Ramaswamy noted the partnership is measured by the “depth of innovation” created for customers.

“Snowflake should have an easier time selling its platform when the LLM [large language model (LLM)] part is backed by a top frontier AI vendor who will work with them on integration and make sure they have early access to new models,” Torsten Volk, an analyst at Omdia told AI Business.

“In return, Anthropic receives access to Snowflake’s large customer base and receives guaranteed revenue at the same time.”

Appian embeds AI Agents in Workflows

Appian has released two major platform capabilities targeting enterprise process automation with built-in governance controls. Agent Studio embeds AI agents directly into workflows with auditability features. Appian Composer accelerates application modernisation from design concept to deployment.

The expansion directly addresses documented enterprise frustration with standalone AI implementations. Appian positions workflow-embedded agents as an alternative to conversational chatbots, citing research suggesting standalone AI approaches fail in enterprise operations environments.

“Many organisations deployed ineffective and expensive, stand-alone AI chatbots in their back-office operations teams,” said Michael Beckley, CTO, Appian.

“Research from MIT shows that approach fails 95% of the time because AI on its own is easily confused by different data contexts. Appian takes a fundamentally different path. We embed specialised AI Agents directly inside operations workflows where they deliver reliable results at massive scale, enabling real-world outcomes, like accurately processing tens of millions insurance quotes per year for one customer.”

Agent Studio enables business users to define automation goals using natural language. Embedded agents determine execution paths within defined constraints, inheriting data access, process context, and governance controls from the platform.

Unlike conversational interfaces, agents operate within workflow boundaries, enabling auditability and compliance controls. Appian claims this approach makes governance “easy,” though specific audit mechanisms remain unspecified in the announcement. The platform reports Agent Studio participants in beta testing found the tool “intuitive or very intuitive.”

Appian Composer generates working applications from requirements specifications and data models. The vendor reports 130 organisations have built over 1,300 applications since the tool’s April 2025 preview.

Composer creates interactive plans for user stories, data structures, processes, and interfaces. The tool provides collaborative workspace for business and IT teams. Appian claims users with “any level of development expertise” can build applications, though application quality, governance, and security review mechanisms remain unspecified.

The data fabric now processes up to 50 million rows with 5x faster write throughput. Enhanced encryption supports transparent data security without performance degradation. This capability targets organisations managing large datasets across compliance-regulated environments.

https://www.appian.com

Iron Mountain buys ACT Logistics

Global information management provider Iron Mountain has acquired ACT Logistics. The acquisition expands the firm's investment in Asset Lifecycle Management (ALM).

The move aims to create a single partner for secure and sustainable IT asset management. This reflects a broader industry trend of combining physical data management with secure IT hardware disposition.

The merger focuses on critical governance and compliance areas. This includes secure data destruction and managing hardware sustainability. The combined services aim to support enterprise circularity goals.

According to Iron Mountain, the company is" investing heavily in Asset Lifecycle Management (ALM) to help organisations maximise asset value, improve sustainability, and manage risk across the technology lifecycle."

Analyst firm Future Market Insights estimates Australia's IT Asset Disposition (ITAD) market is projected to grow at a compound annual rate of 11.2% through 2032, fuelled by strict e-waste disposal regulations and the increasing demand for secure and sustainable IT asset management.

https://www.ironmountain.com

archTIS adds US Data Security Firm

Australian data-centric security provider archTIS has completed its acquisition of US-based Spirion, a provider of sensitive data discovery and classification solutions.

The deal brings 150 enterprise customers and 38 employees to archTIS, expanding its North American footprint across healthcare, financial services, government and education sectors. Financial terms of the acquisition were not disclosed.

The acquisition merges Spirion's data discovery and classification capabilities with archTIS' Attribute-Based Access Control (ABAC) technologies.

The combined portfolio addresses data security and governance across cloud and on-premises environments, including Microsoft 365 and SharePoint.

"The integration of Spirion's data discovery platform with our secure collaboration and access control technologies creates one of the most comprehensive Zero Trust data-centric security portfolios available today," archTIS CEO Daniel Lai said.

Kevin Coppins, Spirion's CEO, has been appointed archTIS Executive Vice President, Commercial Enterprise Solutions and General Manager of the Americas. Ryan Tully, Spirion's Chief Product Officer, continues in that role for the combined entity.

Coppins said the combined solutions "create the enterprise data control point, the critical juncture where data-related decisions are made, policies are enforced, and workflows for data security, compliance, and usage converge".

The acquisition positions archTIS to address growing demand for zero trust architecture and data-centric security solutions.

Organisations face increasing pressure to discover, classify and govern sensitive data across hybrid environments while meeting complex compliance requirements.

The combination of DSPM and Data Access Governance (DAG) capabilities enables organisations to identify where sensitive data resides, classify it appropriately, and enforce granular access controls - addressing a critical gap in many enterprise security programs.

RecordPoint acquires Redactive AI

Data governance vendor RecordPoint has acquired Redactive AI, an Australian artificial intelligence startup.

The move addresses rising data security concerns amid rapid enterprise AI adoption. It combines RecordPoint's governance platform with Redactive's tools for sensitive data discovery and classification.

The acquisition targets highly regulated industries. RecordPoint's clients include Westpac, NAB, APRA, and ASIC.

Redactive, founded 18 months ago, secured HESTA and PEXA as flagship clients. Financial terms of the deal were not disclosed.

The announcement noted Redactive's founders "will soon announce a new AI venture".

RecordPoint CEO Anthony Woodward said Redactive's solution will enhance AI security.

"With over 8 million pieces of data moving through our platform per day, RecordPoint is now handling more data transactions than the NASDAQ," said Woodward.

"Redactive's solution will be instrumental as we look to bolster our capabilities with AI security and governance and continue to scale our operations.

"This is to ensure that, as the volume of data we handle for our clients exponentially grows, our capability grows with it."

"The speed and precision with which the Redactive platform was built made it an ideal candidate for our plans for future growth. While there's still room to grow in both the US and Australia, we're eager to bring RecordPoint's solution into new markets in Asia and Europe."

Redactive co-founder Alex Valente said the deal scales Redactive's impact globally. He called it "one of the first acquisitions of an AI-powered enterprise solution in Australia".

https://www.recordpoint.com

Microsoft Embeds AI Defence in M365

Microsoft is including its Security Copilot AI platform in all Microsoft 365 E5 subscriptions at no additional cost. T

The move is designed to address the global cybersecurity workforce shortage, which has reached four million unfilled positions according to the World Economic Forum. Microsoft Corporate Vice President Dorothy Li stated the challenge is "not just to keep pace - but to leap ahead."

Security Copilot delivers AI-powered agents embedded directly into Microsoft Defender, Microsoft Entra, Microsoft Intune and Microsoft Purview. These agents accelerate investigations, automate routine tasks and help security teams shift from reactive to proactive strategies.

Microsoft 365 E5 customers receive 400 Security Compute Units per month for every 1,000 user licences, up to 10,000 SCUs monthly. This allocation is designed to support typical enterprise scenarios without extra spending.

The company is expanding its agent portfolio with more than 40 new offerings. This includes 12 Microsoft-built agents and over 30 partner-built agents available through the Microsoft Security Store.

New Microsoft agents cover security operations, identity and access management, endpoint management and data security. Identity administrators can deploy agents to remediate risky users, optimise Conditional Access policies and streamline access reviews.

Data security professionals can use agents in Microsoft Purview to discover, analyse and remediate sensitive data risks. IT administrators gain access to agents that convert requirements into policies and assess changes before implementation.

Li said Security Operations Centre analysts have detected malicious emails up to 550 per cent faster using the Phishing Triage Agent. Identity administrators achieved up to 204 per cent greater accuracy identifying missing Zero Trust policies with the Conditional Access Optimisation Agent.

The interactive agent experience, now in public preview, enables focused conversations with each agent. Security Copilot draws on Microsoft's threat intelligence from more than 100 trillion daily signals and integrates with Microsoft Sentinel.

A new enterprise knowledge integration feature, currently in preview, allows agents to access organisational data. This delivers contextual recommendations specific to each environment.

Customers can also build custom agents using Security Copilot's Agent Builder and APIs. Since September 2025, customers have created more than 370 unique agents tailored to their specific requirements.

Datadog Tracks Cloud Outages

Observability platform Datadog has released a free public tool that monitors the health status of more than 30 SaaS providers and 13 AWS services using aggregated customer data and AI models.

Updog.ai provides real-time visibility into service degradations for platforms including Amazon, OpenAI, GitHub, Slack, Stripe, ServiceNow, Zendesk and Zoom. The tool is available to anyone without requiring a Datadog account.

The service uses anonymised Application Performance Monitoring telemetry data from thousands of organisations, analysed through a Bayesian model that identifies abnormal error rates across independent customer environments. Datadog claims this approach can detect issues faster than vendor-controlled status pages.

The tool provides up to 90 days of historical degradation data, enabling IT teams to identify recurring reliability issues that may affect business operations or customer-facing services.

For organisations managing complex digital infrastructure and compliance requirements, the service offers independent verification of vendor-reported service status. This addresses a longstanding challenge where IT teams must rely solely on provider-controlled status pages during incidents.

The release comes as organisations increasingly depend on multiple SaaS providers for critical business functions. Outages or degradations can cascade across interconnected services, making independent monitoring tools valuable for risk management and business continuity planning.

Datadog describes Updog.ai as the first step in a broader initiative, according to the vendor. Planned features include GPU availability monitoring for AI infrastructure teams, spot interruption monitoring for cloud workloads, and cyber attack vector tracking.

While several third-party services aggregate vendor status pages, Datadog claims its approach using customer telemetry data provides earlier detection.

https://updog.ai

AvePoint targets SaaS sprawl

AvePoint has expanded its Confidence Platform with data protection for five additional SaaS applications and new governance capabilities for Microsoft Copilot Studio agents.

The platform now provides automated backup for Monday.com, Docusign, Smartsheet, Okta and Confluence. It also supports Google GCP Virtual Machines, extending multi-cloud backup capabilities.

The update addresses growing complexity as organisations use an average of 275 SaaS applications, according to a study cited by the vendor. Separately, 92 per cent of organisations operate in multi-cloud environments.

The platform now provides visibility into Copilot Studio agent lifecycles, including where agents originate and their compliance status. This allows organisations to monitor what data agents can access across SaaS platforms.

"Agentic AI is helping organisations streamline workflows and realise productivity gains, but agents often have access to sensitive business data," said John Hodges, Chief Product Officer at AvePoint. "Without transparency into how agents are deployed, companies run the risk of data breach and compliance gaps."

The feature aims to help teams govern agent activity and enforce security policies as AI agent deployment expands.

AvePoint also launched an Operational Efficiency Command Center, the fifth command center in its platform. The dashboard surfaces governance metrics including policy violation trends, remediation speed, workspace lifecycle status and service request volumes.

The tool is designed to help compliance and governance teams quantify the impact of data governance practices and demonstrate return on investment.

John Peluso, Chief Technology Officer at AvePoint, said the platform extends across multi-cloud environments to provide centralised governance and protection.

https://www.avepoint.com/au

Automation Anywhere Acquires Aisera

Automation Anywhere has acquired conversational AI specialist Aisera in a strategic move designed to redefine IT Service Management for the AI era and accelerate enterprise automation.

The acquisition creates what Automation Anywhere describes as "the most comprehensive agentic automation portfolio" by combining its Agentic Process Automation (APA) with Aisera's self-service AI agents for IT, HR and customer service.

The combined entity aims to deliver autonomous operations across departments while reducing organisations' reliance on traditional seat-based ITSM products by up to 40%.

"With this acquisition, we're expanding the value we offer to customers, broadening our market reach, and accelerating the momentum of our GenAI business - already the cornerstone of our growth," said Mihir Shukla, CEO and Co-founder at Automation Anywhere.

The companies share a vision for what they call the "Autonomous Enterprise," where up to 80% of work is fully automated or assisted by AI agents, with human work elevated to higher-value tasks.

Through the merger, organisations gain access to an integrated technology stack for enterprise automation, including conversational AI agents with domain-specific training, Agentic Process Automation systems, a Process Reasoning Engine, and the Mozart Orchestrator for managing complex processes across platforms.

Aisera's team of more than 100 AI engineers will join Automation Anywhere as part of the deal, though financial terms were not disclosed.

Abhi Maheshwari, CEO of Aisera, called the merger "an agentic automation powerhouse" that will "redefine the future of enterprise work and deliver exponential value for businesses across the globe."

Adeptia has rebranded its Connect platform as Adeptia Automate, adding AI-powered data mapping and enhanced deployment options aimed at reducing manual integration work for enterprise IT teams.

The Chicago-based company's updated platform addresses what it calls "First-Mile Data" challenges - the external data entering organisations that often requires manual formatting and mapping before it can be processed by internal systems. The rebrand introduces an AI mapping co-pilot designed to automate data transformation tasks without custom coding.

New features include support for Amazon EKS and OpenShift ARO deployments, alongside existing Azure EKS and on-premises options. The platform has added ISO/IEC 27001 compliance capabilities, which Adeptia claims will appeal to organisations managing regulated data.

https://www.adeptia.com

NCS and Newgen team on low-code

Technology services firm NCS Australia has announced a partnership with low-code platform provider Newgen Software to deliver digital transformation solutions.

The partnership combines NCS Australia's consulting capabilities with Newgen's NewgenONE platform, targeting organisations in finance, insurance, public sector and utilities seeking to modernise legacy systems.

Glenn Irvine, National Lead for Partner Solutions at NCS Australia, said the partnership aims to help clients adapt existing technology infrastructure without major IT overhauls. "Our approach caters specifically to enterprise clients, with experience in domains such as superannuation, insurance, and banking, providing the skill set necessary to connect and enhance legacy systems," said Irvine.

Krishna Kumar, ANZ Country Head at Newgen Software, said: "Our partnership with NCS Australia will enable organisations across sectors to overcome technical debt and reimagine their operations using a robust and future-ready platform.

"I'm excited to see this partnership help both of our teams drive innovation and efficiency, helping enterprise-scale businesses be more future-ready and resilient."

The announcement comes as Australian enterprises increasingly seek low-code solutions to address digital transformation challenges while managing complex compliance requirements.

NCS Australia, a Singtel Group subsidiary, provides technology services across 57 specialisations to government and enterprise clients. The company has 13,000 staff across Asia Pacific.

Newgen Software was recognised in Forrester Research's Low-Code Platforms for Professional Developers Landscape Report (Q4 2024) for AI agents, business intelligence and integration capabilities.

https://www.ncs.co https://www.newgensoft.com

Archive360 adds Microsoft AI

Archive360 has partnered with Microsoft to deliver AI-powered eDiscovery and compliance investigation capabilities using Azure OpenAI in Foundry Models.

The collaboration integrates Archive360's data archiving platform with Microsoft's AI technology to enable compliance officers and investigators to detect and analyse potential policy violations across archived communications.

The centrepiece is Archive360's AI Discovery Investigator feature, which allows users to initiate investigations using natural language prompts. The system scans archived emails, Microsoft Teams messages and other digital communications to identify potential misconduct.

When potential policy violations are detected, the platform automatically creates eDiscovery cases and applies legal holds to relevant data. This addresses challenges faced by compliance officers, HR investigators and insider threat analysts who must analyse large volumes of archived data.

"By combining our governed data platform with Microsoft's AI capabilities, we can enable organisations to conduct thorough investigations while maintaining the strict data controls that regulated industries require," said Dan Manners, Vice President of Product Strategy at Archive360.

The integration uses Azure's ecosystem to unify data from multiple enterprise systems and communication platforms. The platform supports both structured and unstructured data, enabling AI to identify patterns across diverse sources.

The system respects granular permission controls and data segregation requirements, ensuring AI analysis only accesses data users are authorised to view. This maintains confidentiality and privacy requirements for sensitive information.

Atturra takes on TCG Process

Australian IT services firm Atturra is expanding its automation portfolio through a new partnership with TCG Process. The agreement sees Atturra become a strategic reseller and services partner for TCG's intelligent document processing (IDP) and process automation platforms.

Atturra will deploy TCG's OCTO platform, which combines process automation with AI orchestration. The integrator will also offer DocProStar, TCG's solution for intelligent document processing.

Atturra claims the "no-code" nature of the platform allows clients with limited technical experience to manage digital transformation projects directly.

The partnership aims to address the shortage of skilled workers by automating labour-intensive data entry and classification tasks.

A specific focus is the "Email Triage Accelerator," designed to automate the management and classification of inbound communication.

"We are excited to welcome Atturra to our global partner network and to help Australian public sector and private enterprises accelerate their digital business initiatives with the power of the TCG Process platform," said Jeff Leibovici, Vice President of Sales, TCG Process, APAC.

"Atturra is an ideal strategic fit for TCG Process with a tremendous footprint among Australian local councils as well as in the education and utilities sectors which can now benefit from our no code platform and reinforce our ability to service the Australian market.

https://atturra.com/au-en/

Drata Launches APAC GRC Hub

Governance, risk and compliance (GRC) software vendor Drata has established a Sydney base to further its push into assisting local organisations with data sovereignty requirements and regulatory complexity.

Daniel Ettenhofer has been appointed as Regional Vice President of Sales for APAC. The company has opened a Sydney office and launched an APAC data centre earlier this year.

The Drata platform automates compliance across more than 25 frameworks, including ISO 27001 and PCI DSS and Australia's Essential 8 framework.

The vendor claims more than 8,000 customers across 60-plus countries, including over 550 in APAC, with the majority of these based in Australia or New Zealand.

The platform automates governance, risk and compliance processes. Drata claims this reduces time spent on preparing and doing annual audits and streamlines security reviews.

"I feel there is a huge opportunity for Drata in the Australian market, given 80% of businesses are under a 2000 employee count. SMB has really been our sweet spot, but we're seeing more and more demand for Drata coming from the Mid-Market and Enterprise.

"There's a huge amount of opportunity in those organisations that where typically people are wearing many hats, and they don't have the ability to effectively manage and scale compliance within their business," said Ettenhofer.

"Typically, compliance becomes an afterthought or a task that's forgotten about, and that introduces gaps and risks into the security posture of the organisation.

"CEOs and their boards have now recognized that GRC becomes a strengthening mechanism for the business posture, for business progression. So, there's a huge amount of momentum in the market.

"We've done studies of organizations that have up to about a thousand people, and they're spending just over a thousand hours per year on compliance-related activity just to pass something like an ISO27001 audit. That could require two or three FTEs.

"The power of compliance automation with Drata is we not only give those hours back, you can shift that headcount into project and valuable delivery work. The automation and the value we deliver from that is key to our value proposition."

Drata is utilising Generative AI to assist with the task of completing vendor security questionnaires which large companies and government agencies increasingly undertake with prospective suppliers.

"The typical -old way of doing that is through receiving a lengthy questionnaire with maybe 100 to 150 questions. It is very time consuming and takes up a lot of resources to answer those.

"But as we are already collecting all the information about an organisations security policies, how compliant they are against all of these frameworks, by leveraging AI and summarization we can populate those answers.

"What may have taken days or weeks to complete surveys and questionnaires, we're doing it in minutes. So again, more time savings that we're handing back to the security and the GRC team within an organization."

https://drata.com

AI Assistant Tackles Skills Shortage

Barracuda Networks has launched a new AI-powered security assistant that enables faster threat identification and response while reducing the burden on resource-constrained security teams.

The Barracuda Assistant, now integrated into the BarracudaONE cybersecurity platform, uses the company's global threat intelligence network to deliver realtime, data-driven guidance for security professionals across both enterprise and managed service environments.

"Barracuda Assistant empowers users of all skill levels to investigate threats quickly and confidently, even in the face of complex attacks," said Brian Downey, vice president of product management at Barracuda.

According to the company, the assistant streamlines workflows by eliminating disruptive context switching between different security functions, allowing teams to transition between vulnerability assessments and incident reviews through a single interface.

Dave Gruber, principal analyst for Omdia, notes that "The use of generative AI and digital assistants in security operations is delivering measurable improvements across many SecOps functions, including threat investigation and response, threat hunting, operationalising threat intelligence, and more."

The assistant is currently available within the BarracudaONE platform and will soon be accessible through Barracuda XDR, Barracuda SecureEdge and the company's support portal.

The tool comes as organisations face increasingly sophisticated attacks with fewer skilled personnel to address them.

By providing intuitive access to actionable insights and configuration recommendations, Barracuda claims the assistant helps close critical skills gaps while reducing human error.

https://www.barracuda.com

Cohesity Partners with Semperis

Data protection vendor Cohesity has launched on-premises cyber vaulting and identity resilience capabilities as organisations face escalating ransomware threats and data sovereignty requirements.

The company announced multiple security features, including FortKnox Self-managed for isolated data vaults within customer data centres.

The on-premises vault option addresses data sovereignty concerns by using obfuscation technologies that the vendor claims prevent discovery even when primary cluster admin credentials are compromised.

Cohesity also introduced Identity Resilience powered by Semperis to protect Microsoft Active Directory environments. The solution combines Cohesity's immutable storage with Semperis' automated forest recovery capabilities.

Active Directory remains a primary target in cyberattacks due to its central role in enterprise authentication systems. However, Cohesity did not disclose pricing or specific recovery time objectives for the integrated solution.

The vendor plans to add 40 cloud connectors by end of 2025 for compute, container, storage and database services across Amazon Web Services, Microsoft Azure and Google Cloud Platform.

Cohesity also announced NetBackup DirectIO, which allows its Data Cloud platform to serve as immutable storage for NetBackup data sources. The company claims the feature delivers up to 53% cost and storage efficiency savings, though independent verification was not provided.

DataProtect will receive hash-based threat scanning capabilities that Cohesity says deliver near-instant search results for indicators of compromise. The vendor will incorporate Google Threat Intelligence into its threat scanning at no additional cost for Enterprise edition customers.

RecoveryAgent, a cyber recovery orchestration tool with embedded malware scanning and AI capabilities, is now generally available. The tool automates testing, rehearsals and recovery execution.

Cohesity expanded its integration with data security platform Cyera to embed data classification and governance capabilities directly into the Data Cloud platform. This enables customers to identify sensitive data within backups and enforce compliance requirements.

"Data sprawl is being driven by increasing cloud and AI adoption," said Amit Raikar, vice president of Strategic Alliances at Cyera. "This increases security and compliance risks that call for enterprise-wide visibility of all sensitive data."

https://www.cohesity.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

AI Assistants that work as a team

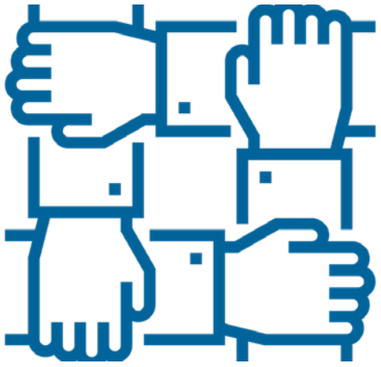

PLANET AI, a German-based artificial intelligence vendor, is expanding its document processing platform with JAIDE, an AI assistant built for organisations handling sensitive data in regulated sectors.

The solution runs on-premises, in private cloud, or hybrid environments, ensuring complete data sovereignty without reliance on US cloud providers.

JAIDE addresses the challenge of extracting actionable intelligence from unstructured corporate data without relying on external cloud infrastructure. The deployment model offers organisations complete control over data handling - a critical consideration for industries bound by stringent privacy regulations.

"JAIDE is not just another chatbot," says Jesper Kleinjohann, CEO of PLANET AI GmbH.

"Our solution uses specialized AI agents that work together. For example, one agent handles and understands incoming documents, another manages business process logic, and a third connects external tools and domain-specific sources. This orchestration delivers far more accurate results than a single generic model."

Instead of forcing users to craft the perfect prompt, JAIDE guides them through pre-built, industry-specific processes.

It also integrates with IDA, PLANET AI's document processing solution

The model-agnostic architecture supports integration of various Large Language Models, adapting to customer requirements and deployment scenarios.

PLANET AI identifies three primary applications: knowledge management for technical service environments (locating manuals and maintenance records), tender analysis for public and private sector procurement, and administrative document creation with AI assistance. Each solution has undergone collaborative development with customers, though the vendor did not name reference accounts or disclose deployment numbers.

https://www.planet-ai.com

Hyland expands cloud integrations

Hyland has announced its Content Innovation Cloud now has native integration for Salesforce, SAP SuccessFactors, Guidewire ClaimCenter, and Workday platforms. The company has also enhanced its document processing tools with generative AI capabilities.

The Salesforce integration allows users to access, import, search and manage documents within their CRM interface.

For insurance sector customers, the Guidewire ClaimCenter integration enables access to documents and data directly within the claims management system. This targets reduced deployment complexity for insurers implementing enterprise content management solutions.

HR departments using SAP SuccessFactors Employee Central or Workday HCM can now access enterprise content without leaving their HR platforms. The Workday integration is approved as a Built on Workday application.

Along with new intelligent integrations, Hyland is enhancing its content intelligence, automation, and management solutions to provide advanced capabilities worldwide.

Hyland's Knowledge Enrichment tool now identifies and masks personally identifiable information for downstream AI and analytics use.

This addresses growing compliance requirements around data privacy as organisations deploy AI systems that process customer and employee records.

The vendor enhanced its Intelligent Document Processing (IDP) platform with what it describes as "agentic" capabilities - AI-driven features that extract contextual information not explicitly stated in documents and derive insights from multiple data points. This automation advancement reflects broader industry movement toward AI agents that can make decisions with minimal human intervention.

For OnBase customers, Hyland introduced Cloud Update Service to automate monthly security and compliance updates. The Nuxeo platform gained support for Google Cloud Platform retention policies and enhanced scalability through decoupled Elasticsearch architecture.

Alfresco users received new performance dashboards and an SAP Information Lifecycle Management connector designed to reduce storage costs and automate governance of SAP content lifecycles.

"We're continuing to meet our customers' evolving needs, helping organisations harness the full potential of their content wherever they are in their digital transformation journey," said Michael Campbell, chief product officer at Hyland.

https://www.hyland.com

InSight DXP targets dark data

Information management provider Iron Mountain has released an AI-enhanced version of its InSight DXP platform.

The update introduces autonomous agents to automate document-intensive workflows and help organisations manage unstructured data.

The platform addresses what the vendor describes as "dark data" - information collected but left unused. New capabilities include agentic AI for workflow orchestration, natural language search, and automated identification of redundant, obsolete and trivial data across enterprise repositories.

InSight DXP now provides a unified view of physical and digital records. Users can track file delivery orders and edit metadata for both formats from a single interface. The system connects to external repositories including ERP and CRM platforms.

For compliance teams, the platform integrates with Iron Mountain's Policy Centre to apply retention schedules and legal holds automatically.

New personally identifiable information redaction features detect and obscure sensitive data without manual intervention.

The update reflects broader industry efforts to extract value from unstructured data, which typically represents the majority of enterprise information assets.

Iron Mountain claims organisations leave up to 80 per cent of valuable information locked in fragmented data.

"We recognise that every enterprise is fighting a costly and invisible organisational drag, where fragmented, dark data inhibits strategic growth and slows decision-making," said Narasimha Goli, chief product and technology officer at Iron Mountain.

The platform's automated discovery and remediation capabilities target ROT data — files that no longer serve business purposes but consume storage and create compliance risks. Organisations can delete unnecessary files or archive content based on the system's governance recommendations.

InSight DXP's search function allows users to query systems using natural language questions. The vendor states the tool provides answers grounded in organisational context, though specific details about the underlying AI models were not disclosed.

Iron Mountain cites regulatory compliance as a key driver for the update. The company claims 23 per cent of banking, financial services and insurance firms identify compliance hurdles as major transformation roadblocks.

The platform is available through AWS, Google Cloud and Microsoft Azure marketplaces.

https://www.ironmountain.com/en-au/services/insight-digital-experience-platform

Microsoft Expands Copilot with Claude

Microsoft has announced the integration of Anthropic's Claude AI models into its Microsoft 365 Copilot suite, marking a significant expansion of its artificial intelligence offerings while maintaining its ongoing partnership with OpenAI.

Enterprise users can now choose Anthropic's Claude Sonnet 4 and Claude Opus 4.1 models in two key areas: the Researcher agent and Copilot Studio. These integrations give organizations greater flexibility to select the right AI for their specific business requirements.

In the Researcher agent, users can now select Anthropic's Claude Opus 4.1 for complex, multistep research tasks. Meanwhile, Copilot Studio now offers both Claude Sonnet 4 and Claude Opus 4.1 as model options for building custom enterprise-grade AI agents.

Charles Lamanna, Microsoft's President of Business and Industry Copilot, emphasized the company's commitment to continuous innovation: "And stay tuned: Anthropic models will bring even more powerful experiences to Microsoft 365 Copilot."

This move represents a strategic shift for Microsoft, which has previously relied primarily on OpenAI's technology for its AI features. The integration comes as Microsoft and OpenAI rework the terms of their partnership and collaborate with other key players across the tech industry.

Microsoft has invested more than $US13 billion in OpenAI and remains its biggest financial backer. A Microsoft spokesperson emphasized that the relationship remains intact, stating: "OpenAI will continue to be our partner on frontier models and we remain committed to our long-term partnership."

According to industry analysts, the decision to incorporate Anthropic's models alongside OpenAI's technology provides significant benefits:

Microsoft leaders have acknowledged Anthropic's Claude models as addressing specific performance gaps observed in recent OpenAI models, especially in tasks demanding high design and output quality.

Before users can access Anthropic's AI models in Microsoft 365 Copilot, administrators must enable them through the Microsoft 365 admin center. These models are hosted outside Microsoft-managed environments.