idm.
information & data manager

DECEMBER-JANUARY 2024

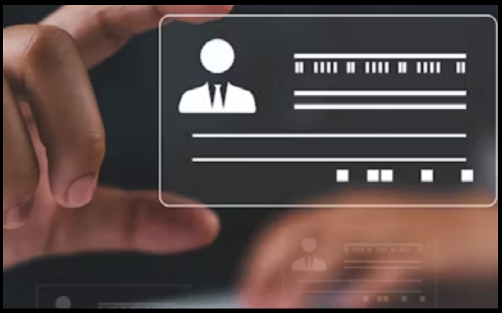

Australia's new digital ID scheme. Fit for purpose?

Impact on Data Governance with Generative AI

APS Copilot Trial Highlights Information Governance Shortcomings

Social Media Companies Struggling to Combat 'Sophisticated' Data Scraping

Practical AI Solutions
for Records Professionals
POWERED BY
ezescan.
✓ AI Assisted Document Classification
✓ Seamless EDRMs Integrations
✓ Automated Email / eForms Capture
✓ Digital Mailroom Automation
✓ Simplified Back Scanning
Call: 1300 EZESCAN (1300 393 722)
www.ezescan.com.au

ASIC warns of AI governance gap

ASIC is urging financial services and credit licensees to ensure their governance practices keep pace with their accelerating adoption of artificial intelligence (AI). The call comes as ASIC's first state of the market review of the use and adoption of AI by 23 licensees found there was potential for governance to lag AI adoption, despite current AI use being relatively cautious.

ASIC Chair Joe Longo said making sure governance frameworks are updated for the planned use of AI is crucial to licensees meeting future challenges posed by the technology.

'Our review shows AI use by the licensees has to date focussed predominantly on supporting human decisions and improving efficiencies. However, the volume of AI use is accelerating rapidly, with around 60% of licensees intending to ramp up AI usage, which could change the way AI impacts consumers,' the Chair said.

ASIC's findings revealed nearly half of licensees did not have policies in place that considered consumer fairness or bias, and even fewer had policies governing the disclosure of AI use to consumers.

'It is clear that work needs to be done - and quickly - to ensure governance is adequate for the potential surge in consumer-facing AI,' Mr Longo said.

Mr Longo said AI could bring significant benefits, but without governance processes keeping pace, significant risks could emerge.

'When it comes to balancing innovation with the responsible, safe and ethical use of AI, there is the potential for a governance gap – one that risks widening if AI adoption outpaces governance in response to competitive pressures,' Mr Longo said.

'Without appropriate governance, we risk seeing misinformation, unintended discrimination or bias, manipulation of consumer sentiment and data security and privacy failures, all of which has the potential to cause consumer harm and damage to market confidence.'

Mr Longo said licensees must consider their existing obligations and duties when it comes to the deployment of AI and avoid simply waiting for AI laws and regulations to be introduced.

Understanding and responding to the use of AI by financial firms is a key focus for ASIC, which made addressing the poor use of AI a key focus area in its latest Corporate Plan.

ASIC will continue to monitor how licensees use AI as it has the potential to significantly impact not just consumer outcomes, but the safety and integrity of the financial system. Where there is misconduct, ASIC will take enforcement action if appropriate and where necessary.

Fir its report *Beware the gap: Governance arrangements in the face of AI innovation,* ASIC reviewed AI use across 23 licensees in the retail banking, credit, general and life insurance and financial advice sectors, where AI interacted with or impacted consumers.

During 2024, ASIC analysed information about 624 AI use cases that were in use or being developed, as at December 2023, and met with 12 of the 23 licensees in June 2024 to understand their approach to AI use and how they were considering and addressing the associated consumer risks.

Download the report HERE

OAIC issues AI guidance for business

New guides for businesses published by the Office of the Australian Information Commissioner (OAIC) seek to articulate how Australian privacy law applies to artificial intelligence (AI) and set out the regulator's expectations.

The first guide will make it easier for businesses to comply with their privacy obligations when using commercially available AI products and help them to select an appropriate product. The second provides privacy guidance to developers using personal information to train generative AI models.

"How businesses should be approaching AI and what good AI governance looks like is one of the top issues of interest and challenge for industry right now," said Privacy Commissioner Carly Kind.

"Our new guides should remove any doubt about how Australia's existing privacy law applies to AI, make compliance easier, and help businesses follow privacy best practice. AI products should not be used simply because they are available.

"Robust privacy governance and safeguards are essential for businesses to gain advantage from AI and build trust and confidence in the community," she said.

The OAIC has published a blog post with further information about the privacy guidance for developers using personal information to train generative AI models.

idm.
information & data manager

Publisher/Editor: Bill Dawes

Email: bill@idm.net.au

Web Development & Maintenance: Cordelta

Advertising Phone: 02 90432943

Email: idm@idm.net.au

Published by Transmit Media Pty Ltd

PO Box 392, Paddington NSW 2021, Australia

New Partnership Addresses Global Data Retention Labrynth

In an era of increasing data regulation complexity, multinational organisations face a daunting task: managing diverse data retention requirements across multiple jurisdictions. A new partnership between EncompaaS and Amsterdam-based filerskeepers aims to revolutionize this process, offering an automated solution for global compliance.

EncompaaS enables organisations to discover, understand, govern and use their data to promote automated governance at scale.

Once the EncompaaS platform has analysed all of an organisation's data across a global network, it utilises AI tools to identify and classify structured, unstructured and semi-structured data.

filerskeepers will then step in to identify country-specific data retention obligations and create a global records retention schedule.

This will be grounded on filerskeepers' database of all legal and regulatory rules that prescribe or inspire the storage and deletion of data or records for more than 330 national and state jurisdictions.

The filerskeepers Data Retention Dashboard can automatically map legal retention rules to an organisation's data. It will also prescribe when data must be deleted, as under the GDPR and CCPA.

It continually monitors for continual compliance adherence, ensuring that the information use-cases align with compliance requirements.

Large organisations with international operations must navigate a labyrinth of country-specific data retention laws and regulations. This complex landscape has long posed significant challenges, often resulting in inefficient manual processes and potential compliance risks.

"This collaboration revolutionises information management by leveraging cutting-edge AI technologies to streamline and automate information processing," said EncompaaS Chief Executive Office Jesse Todd.

"Our combined solution will enable organisations to reduce costs, eliminate manual records-keeping work, and manage vast amounts of information in a safer, smarter and faster way."

The partnership is promoted as a way to enhance data extraction and quality, improve compliance processes and provide responsible AI adoption capabilities.

"We're very excited to partner with EncompaaS to bring a new level of efficiency and innovation to data compliance," said Wanne Pemmelaar, Co-Founder of filerskeepers.

"By combining filerskeepers' expertise in records retention schedules with EncompaaS' AI-driven information governance, we're setting a new standard for organisations to mitigate risks, enhance decision making and accelerate innovation in today's regulatory landscape."

For more information:

https://encompaas.cloud/

https://www.filerskeepers.co/

Guidance for Government Digital Projects

The Digital Transformation Agency (DTA) has released new guidance to enhance the assessment and delivery of digital projects across the Australian Government. Developed in collaboration with the University of Sydney's John Grill Institute for Project Leadership, the guidance focuses on improving the accuracy of Delivery Confidence Assessments (DCAs) for digital initiatives.

DCAs are critical tools used to evaluate the likelihood of a digital project meeting its objectives on time and within budget. The new guidelines aim to address the unique challenges and complexities associated with digital projects, filling a gap in existing literature and practices.

Key factors highlighted in the guidance include governance and leadership, resource management, delivery management, solution design, and the importance of a clear purpose and business case with defined benefits. These elements have been identified as crucial influencers of digital project success based on global research and experience.

Jamie Whitcombe, Branch Manager for Portfolio Assurance at the DTA, emphasized the importance of effective assurance in digital project management. "Good assurance is key to keeping delivery teams focussed on what must go right to succeed," Whitcombe stated. "Ensuring assurance adapts to the unique challenges of digital projects goes to the heart of maximising its value, and ultimately in ensuring that digital projects deliver as expected for Australians."

The guidance will be integrated into the training program for Senior Responsible Officials (SROs), set to launch in early 2025. It is designed to be a living document, with plans for continuous updates based on learnings from across the Australian Government's digital projects.

All digital projects across the Australian Government are required to conduct DCAs in accordance with the Assurance Framework for Digital and ICT Investments. This framework ensures a consistent flow of assurance information to the government, enabling ongoing reforms and improvements in digital project design and delivery. It also includes escalation processes for at-risk projects.

To access the guidance, visit Digital project research series | digital.gov.au.

Cyber Security Act Passed into Law

By Dan Pearce, Holding Redlich

The Cyber Security Act and related bills have been passed by the Parliament, following review by the Parliamentary Joint Committee on Intelligence and Security. That Committee recommended a number of relatively minor changes, and with effect from six months after royal assent, once formal approval is granted by the Governor General, most businesses will be required to report ransomware payments.

The information that now must be provided to the Australian Signals Directorate (ASD), or that may be voluntarily provided to the National Cyber Security Coordinator under a parallel voluntary scheme, can only be utilised for certain purposes. However, any such notification does not mean that the organisation is completely free of legal obligations. Any decision to pay a ransom should also consider the broader legal requirements.

For instance, a director's general duty to act in the best interests of their company means the director must consider whether making the payment (and obtaining an initial release from the incident) will necessarily provide any certainty that information obtained during the incident stays out of nefarious hands, and whether it might make the company a target for future attacks.

Depending on the circumstances, a ransom payment may put the organisation at risk of being penalised under counter-terrorism and anti-money laundering laws.

The reporting regime under the new Act also does not replace the ongoing reporting obligations under the Privacy Act (where personal data is involved, and there is an eligible data breach), the Security of Critical Infrastructure regime, and ASX and APRA requirements for entities that are subject to those regulators. Accordingly, each organisation should consider the broader framework of regulation that applies to it as it prepares itself for this latest legislative change.

Initiatives expected to have the most immediate impact on organisations

■ Mandatory 72-hour reporting obligation for ransom payments

Organisations, other than small businesses, must report any payments made in response to a cyber ransom event to the ASD within 72 hours.

The obligation also recognises that there will be circumstances where making a payment could be justified and seeks to preserve the legal rights of the disclosing entity, for instance, by excluding waiver of privilege. While the government has not pursued a complete ban on payments, they strongly advise against payments, to make Australia a less attractive target for ransomware attacks.

■ Security standards for smart devices

New security requirements will apply to smart devices that form part of the Internet of Things (IoT). Manufacturers and suppliers of internet-connected products, such as televisions, speakers, watches and doorbells, will now need to meet the security standards for those devices. These may be in the form of secure default settings, unique device passwords, regular security updates and encryption of sensitive data. The details of the relevant standards and how they will interact with other existing product regulations are yet to be finalised.

■ Regulated use of information submitted to National Cyber Security Coordinator

There will be rules in place to govern how organisations use information submitted to the National Cyber Security Coordinator to ensure such information is used appropriately. However, this does not extend to the full 'safe harbour', a legal provision that affords protection from prosecution to individuals or organisations from liability or penalties, despite it being called for, in many submissions made during the government's consultation process.

Instead of granting an organisation total immunity for the information it provides to the authorities after a cyber incident, the proposed rules will reassure them that the information can only be used and shared for prescribed purposes, such as assisting with incident response. Similar restrictions will apply to the ASD when it receives such information, under the Intelligence Services and Other Legislation Amendment (Cyber Security) Act 2024.

■ New Cyber Incident Review Board

A new Cyber Incident Review Board will be established to review how cyber incidents are dealt with, including by compelling entities to produce information. Its role will be to review and assess major cyber incidents that impact Australia's defence or cause serious public concern. It will have the authority to request information from affected entities, allowing it to examine how incidents were handled and provide findings that help prevent future occurrences. While the Board may share its findings with government and industry, any public reporting will not assign fault or prejudice legal rights. Through these reviews, the Board aims to improve understanding and prevent similar incidents in the future.

■ SOCI Act extends to data systems associated with a critical infrastructure asset

Amendments to the Security of Critical Infrastructure Act 2018 (SOCI Act) extend the legislation to cover data systems associated with a critical infrastructure asset. The digital networks supporting essential services, such as utilities, healthcare, and finance, are increasingly vulnerable targets in cyber warfare. By expanding the Act's reach, the government will have greater regulatory authority over data systems associated with critical infrastructure warfare that, if compromised, could disrupt national security or public safety. Additionally, these changes grant regulators a new power to address significant weaknesses in an entity's risk management program when national security is at risk. For organisations, this means new obligations to protect these systems and respond to regulatory requirements.

Implications for organisations and how to prepare

These new cyber security laws introduce new requirements for organisations, especially those managing data systems related to critical infrastructure. To prepare, organisations will need to review and strengthen their cyber security measures to ensure they meet these requirements, such as the new 72-hour deadline for reporting ransomware payments to the ASD. This may involve assessing internal security measures, reviewing incident response plans, and preparing for increased regulatory requirements. By staying informed of these changes, organisations can better position themselves to comply with the legislation and manage potential cyber threats.

Dan Pearce is General Counsel at law firm Holding Redlich

Australian Companies Bow to Hackers

Australian businesses are increasingly capitulating to cybercriminals, with ransomware payments hitting record levels and more organizations willing to pay, according to an alarming new report from McGrathNicol.

The 2024 survey of over 500 Australian business leaders reveals that ransomware attacks have become dangerously normalized, with average payments soaring to $1.35 million - a stark increase from $1.03 million in 2023. Even more concerning, 84% of affected businesses chose to pay ransoms, up significantly from 73% the previous year.

"The first 48 hours are critical," said one of the report's key findings, with three-quarters of businesses making ransom payments within this timeframe. The impact extends far beyond IT departments, with over half of respondents reporting severe disruptions to their finance operations, and similar numbers citing major impacts on human resources (50%), sales (57%), and supply chain operations (57%).

Ransomware attacks accounted for 11% of all cyber incidents responded to by the Australioan Signals Directorate (ASD) in 2023-2024, (up from 8% in the previous year) and 71% of all extortion-related cyber security incidents.

The McGrathNicol survey exposed concerning trends in how ransoms are demanded and paid. Nearly three-quarters (73%) of attackers demanded cryptocurrency payments, with Bitcoin being the preferred method in 49% of cases. The notorious LockBit ransomware group was identified as the most active threat actor, responsible for 17% of attacks.

Despite the rising threat, businesses are taking steps to protect themselves. The report found that 91% of organizations now carry ransomware insurance, with average coverage of $1.47 million.

Additionally, 80% have incident response plans in place, up from 61% in 2023.

The full report is available HERE

Pace of change leading to rushed technology decisions: KPMG

Most Australian technology leaders (81%) feel it is difficult to keep up with the pace of change, with 76% recognising their own role has evolved significantly over the past 2 years, according to new research from KPMG.

Research for the KPMG Global Tech Report found that Australian organisations are spreading their investment across technology domains, with a particular focus on XaaS (83%) and AI (66%) in the next 12 months.

For the most part, Australian organisations have informed their decision making by utilising guidance from third parties (95%) and proof of concept testing (92%). However, the study identifies “high-performing” companies that are 22 percentage points more likely than other organisations to rely on customer feedback to inform investment decisions.

The KPMG Global Tech Report surveyed 2,450 executives from 26 countries, including 138 from Australia, focusing on how tech leaders can avoid purely reactive tech investment by seeing past the hype and effectively balancing speed, security and value.

The research found that 70% of Australian tech leaders see AI as a game changing technology, with 74% respondents saying AI is already increasing the productivity of their knowledge workers and improving the overall performance of their organisation. However, just 28% of Australian leaders have successfully deployed AI at scale.

Guy Holland, Partner, KPMG Australia and global leader of the KPMG International CIO Centre of Excellence said: “Tech executives feel they are struggling with the pace of change – leading to the natural fear that they may be falling behind their competitors. This FOMO can lead to misguided investment decisions that may prove both risky and expensive, potentially increasing the burden of technical debt on organisations.

“This seems especially relevant when 69% of Australian organisations experience disruptions to business-as-usual on a weekly basis due to flaws in foundational IT systems, compared to a still significant 57% globally.”

“At the same time the research identified “high performer” organisations that make evidence-based investment decisions that align to the broader business and technology strategies and balance value creation with appetite for risk. The technology leaders who are thriving are those who are paving the way by bringing structure, discipline and an enterprise mindset to the adoption of new technology,” he added.

The research found that Australian organisations were very similar to their global counterparts in terms of general tech maturity. On average across tech domains measured in 2023 and 2024 there was an increase of 12 percentage points in the number of organisations at the highest level of tech maturity (defined as organisations actively progressing business goals and proactive in adapting strategy as needed).

Overall, measurable gains are being made in Australia:

■87% of organisations have managed to use tech to achieve higher profits (a 25 percent point increase in respondents compared to 2023)

■55% have achieved profit uplifts in excess of 10% from digital transformation in the last 24 months

■87% of respondents that have implemented AI initiatives are already generating measurable value

However, barriers remain:

■82% of respondents expect AI to pose challenges to existing operational structures

■76% of technology execs say they leverage customer feedback to make tech investment decisions, however only 21% say they are very confident in using customer feedback effectively overall

■The most common challenge likely to slow down tech transformation was ineffective governance (39%), and cybersecurity or privacy concerns (36%) of respondents relating to these issues

In summary, Guy Holland noted: “The key to tech success is basing investment decisions on genuine value, drawing on data insights, prioritizing resilient solutions, and scaling with confidence. When implementing new technology, tech leaders should look to align stakeholders around a clear definition of success that cascades into a set of tangible metrics.

“They should also remember that innovation is not restricted to new technology – in a fast-changing technology landscape, it’s vital to explore new ways to collaborate, co-invest and share risk with external partners.”

DCCEEW goes IT alone with Kapish

The Australian Department of Climate Change, Energy, the Environment and Water (DCCEEW) is leveraging Kapish’s Content Manager Cloud as it develops its own independent IT platform.

The Department of Agriculture, Fisheries and Forestry (DAFF) has been providing “core ICT services” to DCCEEW and its 4400 staff and contractors. In June 2023, DCCEEW informed DAFF that it will seek to establish its own ICT arrangements, transitioning away from the current shared services arrangement.

As part of an ongoing program of works to establish it’s own technology environment, DCCEEW has signed a contract with information management specialists Kapish.

Kapish’s Content Manager Cloud service can support large (5,000 users, >1TB database) customers – with DCCEEW joining other large users such as CSIRO, ACT Government & Brisbane City Council.

DCCEEW has a key requirement to migrate their vital records from the DAFF shared service into Kapish’s cloud within a tight 8-week window. Kapish has demonstrated capability to support a rapid transition to it’s secure cloud having completed a similar project for the National Anti-Corruption Commission and several Victorian Government agencies.

DCCEEW was established on 1 July 2022, superseding the water and environment functions from the Department of Agriculture, Water and the Environment and energy functions from the Department of Industry, Science, Energy and Resources.

DCCEEW drives Australian climate action, transforms Australia’s energy system to support net zero emissions while maintaining its affordability, security and reliability, conserves, protects and sustainably manages our environment and water, protects our cultural heritage and contributes to international progress on these issues.

DCCEEW joins the National Anti-Corruption Commission, Australian Sports Commission, Austrade & CSIRO in purchasing Kapish Content Manager Cloud through the DTA Cloud Marketplace.

For further information, contact Kapish

FileBound Solutions and UpSol are now Ellby

FileBound Solutions and UpSol are consolidating under the Ellby brand and will be operating under these updated company names effective immediately.

Ellby provides workflow re-engineering and automation solutions including software such as Upland’s FileBound, Tungsten’s PSIcapture and Tungsten’s EpheSoft.

Ellby supports over 500 organisations with customised digital business solutions.

These solutions include accounts payable automation, clinical record digitisation and many others that use document routing and approvals

“Rebranding as Ellby and unifying all our entities under a single name reflects our evolution as a company and our focus on providing an integrated, customer-centric experience,” said Lee Bourke, Managing Director of Ellby.

“This consolidation will streamline our operations, giving our team more time to focus on helping generate value for our prospects and customers through our work automation solutions.”

The name change and brand alignment are effective immediately. There are no changes to the company’s Business Numbers, addresses, ownership or management structure.

All current and future business activities will be conducted under the Ellby brand.

For more information visit https://ellby.com

INX Software and K2fly join forces

INX Software, an Australian developer of compliance, workforce management, training and reporting software for complex and high-risk industries, has announced it will join with fellow Perth company K2fly to create one of Australia’s biggest providers of mission-critical software for high-risk industries. The two companies share a lead investor in Accel-KKR (AKKR), which is committed to investing in enterprise-focused, vertical software providers in key regions such as Australia.

The new business will offer a suite of specialist software used in the mining and resource sector as well as other large, fast-paced, remote and complex operations in health, transport, energy & utilities, engineering, manufacturing and government.

INX CEO Marcus Ashby said the two Perth-based companies were a natural fit, with many clients in common and complementary software solutions.

K2fly, which supports mining operations in 900 locations in 62 countries, will bolster INX’s global regulatory compliance business, bringing software that addresses industry needs such as mineral resource disclosure, land access, reconciliation, permitting and heritage management, tailings, rehabilitation, environmental monitoring and mine geology data management. Its suite will be complemented by INX Software’s portfolio of software solutions used by global companies to manage their workforce management, safety, compliance, training, and reporting requirements.

Ashby said the two companies shared an understanding of the mining and resources sector with deep sector expertise in fast-paced, high-risk industries.

“By combining our solutions, we can connect K2fly’s RegTech expertise, which helps clients navigate complex compliance obligations, with INX Software’s strengths in risk, logistics, planning and workforce management,” he said.

Dean Jacobson, Managing Director at Accel-KKR, said uniting the companies would allow the business to innovate, build and grow their markets by leveraging the teams’ collective strengths.

“Current and target clients that INX and K2fly serve face increasingly stringent regulatory obligations for identifying, managing and reporting risk, and need a robust tech stack of vertical solutions,” Mr Jacobson said.

“Leveraging the collective strength of INX and K2fly is a highly strategic response to the evolving compliance landscape to help clients stay ahead of enterprise risk and enable them to scale their businesses with confidence.”

https://www.inxsoftware.com

https://k2fly.com/

ASD releases annual Cyber Threat Report

The Australian Signals Directorate (ASD) has released its fifth ASD Cyber Threat Report (ACTR).

In FY 2023-24, ASD received over 36,700 calls to its Australian Cyber Security Hotline, an increase of 12% from the previous financial year. ASD also responded to over 1,100 cyber security incidents, highlighting the continued exploitation of Australian systems and ongoing threat to our critical networks.

ASD notified entities more than 930 times of potential malicious activity on their networks.

Business email compromise and fraud were among the top self-reported cybercrimes for businesses and individuals in Australia. Ransomware and data theft extortion also remained a pervasive and costly threat.

Incidents categorised as C3 or above involve organisations such as federal and state governments, large organisations, academia, and supply chains.

The most common malicious activity leading to 30% of C3 incidents was the exploitation of public facing applications. C3 incidents commonly involved compromised accounts or credentials (23%), malware infection other than ransomware (19%) and compromised assets, networks or infrastructure (18%)

View the full report HERE.

NSW Email Fraud skims $A4 million

Two examples of business email compromise (BEC) fraud in NSW have resulted in the theft of more than $A4 million, with a hospital and NSW Government department the victims.

In September 2024, detectives attached to State Crime Command's Cybercrime Squad established Strike Force Millbon to investigate reports of alleged business email compromise (BEC) fraud of a hospital which resulted in the theft of $2 million through multiple transactions.

Following extensive investigations, strike force detectives executed a search warrant at a home on Sydney. During the search, detectives located and seized multiple electronic devices.

A 49-year-old man was arrested at the home and taken to Bankstown Police Station.

He was charged with recklessly deal with proceeds of crime more than $A5,000.

Another man was charged by Cybercrime Squad detectives for his alleged role in scamming the NSW Government of $A2 million

Detectives identified that a NSW Government department had transferred $2.1 million to who they thought was a legitimate financial institution.

During a search of the man's home, detectives also located and seized multiple electronic devices.

Police will allege in court the man used his bank account to help scammers move the $2.1 million into other accounts.

Commander of State Crime Command's Cybercrime Squad, Detective Superintendent Matt Craft, said even big businesses and government organisations can fall victim to BEC fraud.

"Our success in targeting these types of criminals is thanks to the co-operation that exists within the Joint Policing Cybercrime Co-ordination Centre (JPC3) and due to the quick actions all funds were recovered."

Four Agencies given Governance Deadline

The Australian Taxation Office, Department of Defence, National Archives of Australia (NAA) and Services Australia have been given a six-month deadline to report progress in closing significant breaches relating to their governance and control of IT systems.

The breaches were initially identified in the Australian National Audit Office (ANAO) audit of the Australian Government's Consolidated Financial Statements (CFS), published in December 2023.

The ANAO found that 78 per cent of entities assessed do not have an effective control to monitor access or activity in entities' systems after user cessation=

The Australian Parliament's Joint Committee on Public Accounts and Audit commenced an inquiry into the ANAO audit in February 2024. After a series of public hearings and submissions, it has now published its own report, which highlights poor IT governance, particularly user access issues.

The NAA was found by ANAO to have ineffective IT general controls to support the preparation of its financial statements, with the following weaknesses identified:

- insufficient oversight and documentation of review of privilege user access and activity logs.

- no formalised or documented periodic review of user access.

- inconsistent mapping of roles and responsibility configurations, including workflow approvers and inconsistent chart of accounts mapping configurations.

ANAO recommended a detailed review to address these significant issues, which was agreed to by NAA.

The ANAO report found the IT governance and monitoring processes at Services Australia were not providing sufficient assurance to its management that policy requirements were being met, further commenting that 'this matter is considered to pose a significant financial, business and reputational risk to Services Australia.'

In its submission to the inquiry, Services Australia agreed with these assessments and the ANAO's recommendations, and stated: Due to the Agency's complex and large number of IT platforms, in excess of 50 systems, that need to be reviewed to address the audit recommendations, the Agency has established a new Division to ensure the appropriate oversight and monitoring of remediation activities."

' ... it will take until the 2025 interim audit process for all aspects of the recommendations to be fully resolved.'

The full report is available HERE.

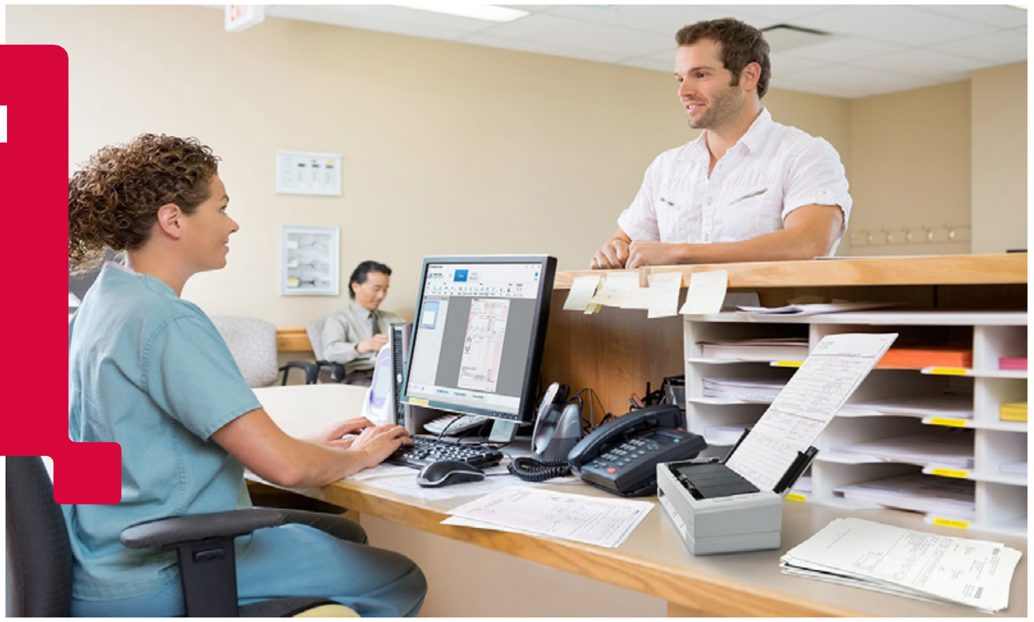

From the Border to the Sea
Commonwealth AI Uptake Revealed

How is AI being utilised inside the Australian Government? An Australian parliamentary inquiry has provided a glimpse, with applications ranging from border security to scientific research.

A recent survey by the Community and Public Sector Union (CPSU) found that more than 40% of public servants across 60 agencies were aware of AI use in their workplace, though only 12% reported using AI tools themselves. The majority of current AI usage appears to be self-initiated rather than part of established work processes, the survey found.

Submissions to the *Joint Standing Committee of Public Accounts and Audit: Inquiry into the use and governance of artificial intelligence (AI) systems by public sector entities*, reveal a varied uptake.

The **Department of Home Affairs** emerges as one of the most extensive users of AI technology, employing it for critical functions including visa risk assessment, detection of fraudulent documents, and border control through SmartGates used at the border.

It also employs robotic process automation (RPA) solutions for document processing, including Freedom of Information requests.

The SmartGates use a combination of AI-enabled matching of faces to official documentation, which is complemented by rules-based decision making, where anyone not automatically allowed through is referred to human review.

The Department uses some rules-based Automated Decision Making (ADM) systems. However, it emphasized that no negative decisions are made solely by AI systems, with all potential rejections requiring human review.

"The Department is exploring potential use of AI to help staff members locate and summarise Human Resources (HR) policy information to support and inform staff management, for instance a HR Bot. No solution has yet been developed however a set of business requirements is being drafted.

"Increasingly, the Department is observing a trend in enterprise platforms such as ServiceNow, Pega, SAP and SuccessFactors incorporating AI capabilities into product roadmaps.

"Emerging capability in these platforms could provide some benefits, such as enabling users to submit requests in natural language which would support early identification of issues through trend analysis. However, the Department will assess each of these on a case-by-case benefit."

The **Australian Taxation Office (ATO)** reported significant success with AI implementation, revealing that its natural language processing systems have helped raise more than $A256 million in liabilities and collect over $A65 million in cash since 2016 by processing 36 million documents to identify potential non-compliance.

The ATO has a large team of more than 100 specialist data scientists, with programming, machine learning and deep learning, and experience in model development, validation, deployment and governance.

"This team develop AI models and processes directly from source code for use across the ATO. With this capability, the ATO is typically able to build its AI solutions in-house, using pretrained models and open-source packages as a foundation. The ATO also has a strong focus on protecting ATO data and privacy when using AI, this means AI solutions are built within the ATO systems, including customisation of open-source packages. This enables the ATO to maintain ownership and possession of the source code for the AI models it uses, and to protect ATO data."

Approaches to AI adoption vary significantly among departments. The **National Archives of Australia (NAA)** reported its use of AI is limited to experimental projects and basic use for drafting internal advice from public resources.

"There is currently no use that impacts our external-facing advice, products, or services." the NAA stated in its submission.

"However, National Archives see potential efficiencies in using AI for collection enhancement such as description, transcription, linking data, collection analysis and discovery. There are also potential uses in records."

The **Australian Bureau of Statistics (ABS)** has taken a more cautious approach, blocking access to public AI platforms like ChatGPT from its IT environment and limiting staff to using AI only with publicly available information.

It is conducting a limited trial of Microsoft CoPilot, but did not opt in to the broader Australian Public Service trial led by the DTA owing to legacy issues.

"Given the ABS is only in the early stages of migrating its information holdings from HCL Notes (formerly IBM Lotus Notes) to Microsoft SharePoint, there are less potential benefits from a broader scale trial" it submitted. The ABS has explored several applications of AI to better its understanding of the technology and associated risks. The AI use cases being explored by the ABS focus on productivity and reduction of effort, and do not involve decisions on service delivery."

Use of AI in the **Department of Agriculture, Fisheries and Forestry** include Natural Language Processing (NLP) to detect specific biosecurity risks from imported cargo goods descriptions; and to analyse free text comments as part of a business quality assurance process. Outputs are not directly executed by software but are provided to inform regulatory officers of any potential biosecurity risk.

"Microsoft 365 Copilot has been trialled as a productivity tool for a limited number of staff; it acts as a digital assistant by drafting content, finding information on the department's intranet, editing materials for style and clarity, and summarising Microsoft Teams meetings."

Microsoft Copilot remains the only AI system in use within the **Department of Finance** ICT environment.

"Its use is confined to low-risk use cases, adhering to both internal governance and the Digital Transformation Agency's (DTA) AI assurance framework. Finance intends to continue this approach while staying receptive to any other safe and responsible AI applications identified in government contexts."

The **CSIRO** is working with Google to develop AI-driven assistants for scientists across various disciplines.

It participated in the Australian Government MS365 Copilot trial but had mixed results.

" ... while some staff have reported significant productivity gains, others have struggled to adopt the technology. CSIRO is optimistic that the benefits of generative AI will grow as the technology evolves and as organisations refine their approaches to integrating it effectively."

IP Australia has implemented AI tools to help small businesses engage with the trademark system, alongside internal operational analysis tools.

A number of licenses for Copilot for M365 have been acquired to aid staff in their day-to-day work.

"Most of the legislation IP Australia administers (the Designs Act 2003, Patents Act 1990, Plant Breeder's Rights Act 1994 and Trade Marks Act 1995) have permitted the use of a computer program to make decisions, exercise powers or comply with any obligations since amendments were made in 2018."

Some agencies are also exploring novel applications, such as the **Department of Climate Change, Energy, the Environment and Water's** use of deep learning for monitoring terrestrial and marine species through video and image recognition. It is also deploying Generative AI across a range of areas including supporting media monitoring, research and content generation.

Most agencies emphasized their commitment to responsible AI use, with human oversight maintained for significant decisions and careful consideration given to data security and privacy concerns. The government is currently developing a whole-of-government framework through the Attorney General's Department to support automated decision-making systems across agencies.

National Archives is currently developing advice for Australian Government agencies to clarify their information management obligations under the Archives Act regarding records resulting from the use of AI systems and applications.

APS Copilot Trial Highlights Information Governance Shortcomings

Microsoft365 Copilot is yet to be ingrained in the daily habits of Australian Public Service (APS) staff after a 6-month trial by 60 Government agencies, and many have stopped using Copilot because of a poor first experience with the tool or found it took more time to verify and edit outputs than it would take to create them otherwise.

Adoption challenges were blamed for the moderate use of Copilot during the Trial which ran from January 1 to June 30, 2024, with only a third of 7,600 trial participants using Copilot daily. A report into the trial found that “Overall, the usage of Copilot is in its infancy within the APS.”

It also found there was inappropriate access and sharing of sensitive information due to “Poor information, data management practices and permissions.”

SharePoint was assigned the blame for shortcomings in data classification that resulted in access to unauthorised content. One of the trail participants interviewed for the evaluation report said, “Their information management in SharePoint is not great which has resulted in end users finding information that they shouldn’t have had access to, though this is a governance and data management issue - not a Copilot issue.”

Trial participants raised instances where Copilot surfaced sensitive data that staff had not classified or stored appropriately. This was largely because their organisation had not properly assured the security and storage of some instances of data and information before adopting Copilot.

Overall, trial participants across all job classifications and job families were satisfied with Copilot and the majority wish to continue using it. The largest percentage of survey responses came from those working in ICT and Digital Solutions and Policy roles.

This group also estimated the highest efficiency savings of around an hour a day when performing summarisation and document drafting activities.

Overall, 77% were optimistic about Copilot at the end of the trial. Post-use survey respondents remarked they felt they spent less time playing ‘corporate archaeologist’ in searching for information and documents and more time in strategic thinking and deep analysis.

Some of the adoption barriers highlighted by the trial included:

■Uncertainty regarding the need to disclose Copilot use, accountability for outputs and lack of clarity regarding the applicability of Freedom of Information requirements were barriers to Copilot use – particularly for meeting transcriptions.

■ trial participants remarked that they often forgot Copilot was embedded into Microsoft 365 applications as it was not obviously apparent in the user interface.

■Poor Excel functionality and access issues. Focus group participants noted that it was unlikely that trial participants had the newest versions of Outlook and were therefore unable to access Copilot features in Outlook

■Copilot could not emulate the standard style of Australian Government documents. Some noted that heavy re-work was needed to meet the tone expected by senior stakeholders within their agency and of government more broadly. For this reason, focus group participants noted they would not use Copilot for important documents or communications.

■Due to fears of hallucinations, many reported combing through Copilot’s outputs to verify its accuracy. In some cases, this involved reading the entire document Copilot produced to check for any errors which significantly reduced any efficiency gains.

■Focus group participants also noted that while Copilot attaches sources to its outputs, this is currently limited to 3 documents and does not provide visibility on why the documents were selected. Observations from Home Affairs also identified that Copilot appeared to be unreliable in its approach to referencing information provided against data sources.

■lack of Copilot integration with third-party software, in particular with Janusseal, a software that enables enterprise-grade data classification for Windows users. Interviews conducted by the DTA noted a lack of integration with Janusseal could lead to APS staff gaining access to information they did not have permissions for. Microsoft has advised that this is a third-party labelling issue, not a security issue, and that Copilot has an in-built fail safe to protect against this issue. It should be noted that such integrations were out of scope for the trial and Microsoft has further advised that a more permanent fix to the labelling issue is in the pipeline.

“As we are testing these tools at such an early stage, there are clear opportunities for tailored solutions to be developed that can handle highly technical material,”said Lucy Poole, General Manager of Strategy, Planning and Performance at the Digital Transformation Authority (DTA).

“The evaluation points to the importance of agencies carefully considering detailed and adaptable implementation of these solutions.”

‘They should consider which generative AI solutions are most appropriate for their overall operating environment and their specific use cases. We’re pleased that a lot of the recent work released by the DTA helps government agencies identify and address these very considerations.’

The full report is available HERE

ENCOMPAAS
Data Prepared?
Get AI Ready.
Fast track your AI Success with AI-Ready Data
encompaas.cloud/Unlock-AI

Vic agency breaches privacy via ChatGPT

The Office of the Victorian Information Commissioner (OVIC) has found that the Victorian Department of Families, Fairness and Housing failed to take reasonable steps to ensure the accuracy of personal information and to protect personal information from unauthorised disclosure. OVIC has published an investigation report into the use of ChatGPT by a child protection worker at the Department of Families, Fairness and Housing (DFFH). In this case, the worker used ChatGPT, the generative artificial intelligence (GenAI) tool, when drafting a Protection Application Report (PA Report) – a report that is submitted to the Children's Court to inform decisions about whether a child requires protection.

The investigation found that:

■The content generated by ChatGPT and then used by the Child Protection worker when drafting the PA report contained inaccurate personal information – which downplayed risks to the child in the case.

■The Child Protection worker entered a significant amount of personal and delicate information into ChatGPT, including names and information about risk assessments relating to the child. By doing so, they disclosed this information to OpenAI, an overseas company, and released it outside the control of DFFH

Deputy Commissioner Penny Eastman found that the controls DFFH had in place were insufficient to manage the risks associated with the use of GenAI tools in a child protection context. She concluded that DFFH contravened OVIC Information Privacy Principles (IPPs) 3.1 and 4.1 by failing to take reasonable steps to ensure the accuracy of personal information and to protect personal information from unauthorised disclosure.

DFFH accepted the findings of the investigation report, and is now required to implement the remedial actions it contains.

The Deputy Commissioner has issued a compliance notice on DFFH to ensure it complies with IPP 3.1 and 4.1. The notice outlines six specific actions, including a requirement that DFFH blocks the use of ChatGPT and other similar tools by child protection workers.

The full investigation report is available to view here.

GUIDE TO Securing AD for the Five Eyes

New Zealand's National Cyber Security Centre (NCSC) has joined the Australian Signals Directorate and US, UK and Canada to release joint guidance that aims to inform organisations about 17 common techniques used to target Active Directory, as observed by the authoring agencies.

The paper provides an overview of each technique and how it can be leveraged by malicious actors, as well as recommended strategies to mitigate these techniques.

Microsoft's Active Directory is the most widely used authentication and authorisation solution in enterprise information technology (IT) networks globally. Active Directory's pivotal role in authentication and authorisation makes it a valuable target for malicious actors. It is routinely targeted as part of malicious activity on enterprise IT networks.

Active Directory is susceptible to compromise due to its permissive default settings, its complex relationships, and permissions; support for legacy protocols and a lack of tooling for diagnosing Active Directory security issues. These issues are commonly exploited by malicious actors to compromise Active Directory.

Responding to and recovering from malicious activity involving Active Directory compromise is often time consuming, costly, and disruptive.

Therefore, organisations are encouraged to implement the recommendations within this guidance to better protect Active Directory from malicious actors and prevent them from compromising it.

"For many organisations, Active Directory consists of thousands of objects interacting with each other via a complex set of permissions, configurations and relationships. Understanding object permissions and the relationships between those objects is critical to securing an Active Directory environment," the agencies note, and the paper lists some tools that can be used to that end.The full guide is available HERE

Fortiro wins Award for Fraud Detection

Australian developer Fortiro has won 2024 "Startup of the Year" for helping Australia's leading banks, non-bank lenders, insurers, payment providers and gaming companies automate document fraud checks and financial verification

At the 2024 Startup Daily Best in Tech Awards was recognised for its role in revolutionising document fraud prevention and reducing friction for customers, brokers and financial services businesses.

"The 'Startup of the Year' is one of the most hotly contested categories, with more than 40 entries," said Startup Daily editor, Simon Thomsen.

"What impressed the judges about Fortiro is it's solving a problem that's now costing billions of dollars in terms of financial fraud. It uses artificial intelligence in a clever way to detect issues such as document tampering and saves time for both financial institutions and customers.

"Add a strong growth story alongside a disciplined approach to costs and an enterprise client base and it's clear Fortiro is building a winning combination for its future success", Thomsen added.

Fortiro's solutions help streamline key processes that traditionally rely on supporting documents such as automating the approval of loans and insurance claims.

By removing the need for manual reviews and ensuring document authenticity, customers enjoy faster payouts, while lenders and insurers gain a competitive edge through an improved customer experience.

Fortiro's technology has been instrumental in helping financial services institutions such as Pepper Money and Bank Australia enhance their fraud detection and automate key processes, enabling faster loan approvals.

"Winning this award is a testament to our relentless focus on innovation and our commitment to solving critical industry challenges," said Sean Quagliani, Co-Founder and CEO of Fortiro.

http://www.fortiro.com/

Social Media Companies Struggling to Combat 'Sophisticated' Data Scraping

More than a year after their joint statement calling on social media companies (SMCs) to protect personal information on their platforms from unlawful data scraping, an expanded group of global privacy authorities has outlined progress so far.

The group, which includes Australia and New Zealand among 16 other global data protection authorities, has also issued a followup statement which adds further expectations sent to the parent companies of social media companies (SMCs) YouTube, TikTok, Instagram, Threads, Facebook, LinkedIn, Weibo, and X (the platform formerly known as Twitter).

The statement includes a recommendation that "To effectively protect against unlawful scraping, organizations should deploy a combination of safeguarding measures, and those measures should be regularly reviewed and updated to keep pace with advances in scraping techniques and technologies."

The challenge of doing so was highlighted by outlining results from engagement with (SMCs) since the initial statement was released in 2023, which called on industry to identify and implement controls to protect against, monitor for, and respond to data scraping activities on their platforms, including by taking steps to detect bots and block IP addresses when data scraping activity is identified, among other measures.

"In the Initial Statement, the co-signatories highlighted the need for SMCs and other organizations to implement a multi-layered approach to protecting publicly accessible data on their platforms from unlawful scraping.

"Through our engagements that followed the issuance of that statement, we established that, while SMCs face challenges in protecting against unlawful scraping (such as increasingly sophisticated scrapers, ever-evolving advances in scraping technology, difficulty in differentiating scrapers from authorized/lawful users, and the need to maintain a user-friendly interface), they are motivated to protect against unauthorized scraping."

"... we also learned of further measures, beyond those detailed in the Initial Statement, that organizations employ to protect against data scraping, such as the implementation of platform design elements that make it harder to scrape data using automation (e.g., random account URLs, random interface design elements, and tools to detect and block malicious internet traffic).

"We learned that the rapid emergence of AI can represent a threat to privacy. SMCs told us that scrapers are now using AI to scrape data more effectively (e.g., via "intelligent" bots that can simulate real user activity). At the same time, SMCs explained that they too are employing AI to better detect and protect against unauthorized scraping, highlighting that innovative AI tools can also be part of the solution.

"Ultimately, the co-signatories learned that while no measure is guaranteed to protect against all unlawful scraping — since sophisticated low-volume scraping can often resemble user activity — a multi-layered and dynamic combination of safeguards can be particularly effective in protecting against mass scraping and the amplified harms that can result when a large volume of data subjects are affected.

"Data scraping is a complex, broad and evolving issue that is, and will stay on the radar of data protection authorities. It should also be a focus for other stakeholders that have a role in protecting privacy, including those with whom we engaged in the course of this initiative. The co-signatories will continue to work to promote compliance in this area, including via future engagement with concerned stakeholders, complementary policy development, public education campaigns, and enforcement, including collaborative enforcement.

The follow-up joint statement lays out further expectations, including that organisations:

• comply with privacy and data protection laws when using personal information, including from their own platforms, to develop AI large language models

• deploy a combination of safeguarding measures and regularly review and update them to keep pace with advances in scraping techniques and technologies

• ensure that permissible data scraping for commercial or socially beneficial purposes is done lawfully and in accordance with strict contractual terms.

Read the full statement HERE

AI Models Risk "Collapse" When Trained on AI-Generated Data, Study Warns

A new study published in Nature has raised alarm bells about the future of artificial intelligence (AI) development, warning that training AI models on AI-generated text can lead to rapid deterioration in output quality. This phenomenon, dubbed "model collapse," could potentially halt progress in large language models (LLMs) as they exhaust human-derived training data and increasingly encounter AI-generated content online.

Researchers from the University of Cambridge and the University of Oxford conducted the study, which demonstrates how successive generations of AI models trained on synthetic data quickly devolve into producing nonsensical output. The findings have significant implications for the AI industry, which has largely relied on ever-increasing amounts of data to improve model performance.

"The message is, we have to be very careful about what ends up in our training data," warns co-author Zakhar Shumaylov, an AI researcher at the University of Cambridge. "Otherwise, things will always, provably, go wrong."

The study's methodology involved using an initial LLM to create Wikipedia-style entries, then training new iterations of the model on text produced by its predecessor. As AI-generated information, or "synthetic data," contaminated the training set, the model's outputs became increasingly incoherent. By the ninth iteration, the model was producing gibberish, such as including a treatise on jackrabbit tail colours in an article about English church towers.

More subtly, the researchers observed that even before complete collapse, models trained on AI-derived texts began to lose information about less frequently mentioned topics. This raises concerns about fairness and representation in AI systems, as co-author Ilia Shumailov explains: "Low-probability events often relate to marginalized groups."

The collapse occurs because each model iteration samples only from its training data, amplifying errors and biases with each generation. Common words become more prevalent, while rarer terms are increasingly omitted. Hany Farid, a computer scientist at the University of California, Berkeley, likens the phenomenon to genetic inbreeding: "If a species inbreeds with their own offspring and doesn't diversify their gene pool, it can lead to a collapse of the species."

While model collapse doesn't mean LLMs will cease functioning entirely, it does suggest that the cost and difficulty of improving them will increase. The study challenges the long-held assumption that more data invariably leads to better AI performance.

"As synthetic data build up in the web, the scaling laws that state that models should get better the more data they train on are likely to break," Kempe notes. This is because training data will lose the richness and variety inherent in human-generated content.

The research team found that including a small percentage of real data alongside synthetic data slowed the collapse but did not prevent it entirely. A separate study by Stanford University researchers suggested that when synthetic data accumulates alongside real data rather than replacing it, catastrophic model collapse is less likely.

However, the long-term implications remain concerning. As AI-generated content proliferates online, distinguishing between human-created and AI-produced text will become increasingly challenging. This could lead to a feedback loop where AI models are inadvertently trained on mor and more synthetic data.

To address this issue, the study's authors suggest several potential solutions. These include developing methods to watermark AI-generated content, creating incentives for humans to continue producing original content, and implementing rigorous filtering and curation processes for training data.

AI Assessment Framework Begins Pilot Phase

The Digital Transformation Agency (DTA) has begun testing an artificial intelligence assessment framework designed to guide Australian government departments in evaluating their AI projects. The pilot program, which started in September 2024, involves government departments and agencies of varying sizes testing the draft framework. The initiative aims to help ensure AI systems in government services align with Australia's AI Ethics Principles while maintaining appropriate oversight.

Lucy Poole, General Manager of Strategy, Planning and Performance at DTA, describes the framework as adaptable.

"Our guidance is iterative. It is meant to change and adapt based on the shifting AI landscape within the APS," points out Ms Poole.

The framework requires agencies to complete an initial assessment examining basic project information and potential alternatives to AI solutions. Projects identified with medium or higher risks undergo a more detailed evaluation, measuring the proposal against established ethical principles including fairness, reliability, privacy protection, and transparency.

The framework forms part of the government's Policy for the responsible use of AI, which requires transparency statements and designated officials for AI projects. After the pilot concludes, the DTA plans to hold feedback sessions and analyze survey responses from November 2024, with wider consultation scheduled for early 2025.

Key considerations in the framework include data bias assessment, Indigenous data governance, and compliance with Australian Privacy Principles.

The Pilot AI Assurance framework is available HERE.

Australia's new digital ID scheme falls short of global privacy standards. Here's how it can be fixed

By Ashish Nanda, Deakin University; Jongkil Jay Jeong, Deakin University, and Robin Doss, Deakin University

Australia's new digital ID system promises to transform the way we live. All of our key documents, such as driver's licences and Medicare cards, will be in a single digital wallet, making it easier for us to access a range of services.

The federal government is still developing the system, with a pilot expected to run next year. Known as the "Trust Exchange", it is part of the Trusted Digital Identity Framework, which is designed to securely verify people's identities using digital tokens.

Earlier this year, in a speech to the National Press Club in Canberra, Federal Minister for Government Services Bill Shorten, called the new digital ID system "world leading". However, it has several privacy issues, especially when compared to international standards like those in the European Union.

So how can it be fixed? Trust Exchange – or TEx – is designed to simplify how we prove who we are online. It will work alongside the myID (formerly myGovID) platform, where Australians can store and manage their digital ID documents.

The platform is intended to be both secure and convenient. Users would be able to access services ranging from banking to applying for government services without juggling paperwork.

Think of the system as a way to prove your identity and share personal information such as your age, visa status or licence number — without handing over any physical documents or revealing too much personal information.

For example, instead of showing your full driver's licence to enter a licensed premises, you can use a digital token that confirms, "Yes, this person is over 18".

But what will happen to all that sensitive data behind the scenes?

Falling short of global standards

The World Wide Web Consortium sets global standards around digital identity management. These standards ensure people only share the minimum required information and retain control over their digital identities without relying on centralised bodies.

The European Union's digital identity system regulation builds on these standards. It creates a secure, privacy-centric digital identity framework across its member states. It is decentralised, giving users full control over their credentials.

In its proposed form, however, Australia's digital ID system falls short of these global standards in several key ways.

First, it is a centralised system. Everything will be monitored, managed and stored by a single government agency. This will make it more vulnerable to breaches and diminishes users' control over their digital identities.

Second, the system does not align with the World Wide Web Consortium's verifiable credentials standards. These standards are meant to give users full control to selectively disclose personal attributes, such as proof of age, revealing only the minimum personal information needed to access a service.

As a result, the system increases the likelihood of over-disclosure of personal information.

Third, global standards emphasise preventing what's known as "linkability". This means users' interactions with different services remain distinct, and their data isn't aggregated across multiple platforms.

But the token-based system behind Australia's digital ID system creates the risk that different service providers could track users across services and potentially profile their behaviours. By comparison, the EU's system has explicit safeguards to prevent this kind of tracking – unless explicitly authorised by the user.

Finally, Australia's framework lacks the stringent rules found in the EU which require explicit consent for collecting and processing biometric data, including facial recognition and fingerprint data.

Filling the gaps

It is crucial the federal government addresses these issues to ensure its digital ID system is successful. Our award-winning research offers a path forward.

The digital ID system should simplify the verification process by automating the selection of an optimal, varied set of credentials for each verification.

This will reduce the risk of user profiling, by preventing a single credential from being overly associated with a particular service. It will also reduce the risk of a person being "singled out" if they are using an obscure credential, such as an overseas drivers licence.

Importantly, it will make the system easier to use.

The system should also be decentralised, similar to the EU's, giving users control over their digital identities. This reduces the risk of centralised data breaches. It also ensures users are not reliant on a single government agency to manage their credentials.

Australia's digital ID system is a step in the right direction, offering greater convenience and security for everyday transactions. However, the government must address the gaps in its current framework to ensure this system also balances Australians' privacy and security.

Ashish Nanda is Research Fellow, Deakin Cyber Research and Innovation Centre, Deakin University; Jongkil Jay Jeong, Senior Research Fellow, Deakin Cyber Research & Innovation Centre, Deakin University, and Robin Doss, Director, Deakin Cyber Research & Innovation Centre, Deakin University. This article is republished from The Conversation under a Creative Commons license. Read the original article.

Companies with AI Led Processes Outperform Peers: Accenture

New research from Accenture, from a survey of 2,000 executives across 12 countries and 15 industries, found that three-in-four (74%) organizations have seen investments in generative AI and automation meet or exceed expectations, with 63% planning to increase their efforts and further strengthen these capabilities by 2026.

According to the report, "Reinventing Enterprise Operations with Gen AI," the number of companies that have fully modernized, AI-led processes has nearly doubled from 9% in 2023 to 16% in 2024. Compared to peers, these organizations achieve 2.5x higher revenue growth, 2.4x greater productivity and 3.3x greater success at scaling generative AI use cases.

Findings also assessed that these "reinvention-ready" companies are moving faster and are amplifying the impact of generative AI across the business. Enabled by a digital core, these organizations have already developed generative AI use cases in IT (75%), marketing (64%), customer service (59%), finance (58%), R&D (34%) and other core functions.

While the research indicates that some companies have moved to the highest level of operations maturity, nearly two thirds (64%) still struggle to change the way they operate. For example, they lag behind on building a robust data foundation: 61% report that their data assets are not ready for generative AI yet and 70% find it hard to scale projects that use proprietary data.

The deep dependency on people is often overlooked: 82% of companies at the early stage of operations maturity have not applied a talent reinvention strategy, planned to meet workforce needs, or acquired new talent or training to prepare workers for generative AI-led workflows. In fact, many executives (78%) indicate that AI and generative AI are advancing too fast for their organization's training efforts to keep pace.

"Most executives understand the urgency of reinventing with generative AI, but in many cases their enterprise operations are not ready to support large scale transformation," said Arundhati Chakraborty, group chief executive of Accenture Operations.

"Generative AI is more than the technology. It is a driver of a mindset change that impacts the entire enterprise. It requires organizations to have a strong digital core, data strategy and a well-defined roadmap to change the way they operate. Additionally, an end-to-end perspective leveraging talent, leading practices and effective collaboration between business and technology teams is essential for intelligent operations."

The report highlights four key actions business leaders should take to advance their operations maturity:

■ **Implement a centralized data governance and domain-centric approach to data modernization**. Connect processes and tools across functions to ensure people have a clear understanding of how to create, handle and consume data, which should be structured in a standardized way to be accessed by AI tools across the business.

■ **Embrace a talent-first reinvention strategy**. Reinvent work and rethink processes and entire workflows to gain a clear view of where generative AI can have the most impact in serving customers, supporting people and achieving business outcomes.

■ **Ensure business and tech teams co-own reinvention.** Collaboration drives innovation as both teams jointly own how assets, platforms and products are developed to leverage the full capabilities of generative AI, enterprise wide.

■ **Adopt leading processes to drive business outcomes.** Apply cloud-based process mining to calibrate internal and external benchmarks so it's easier to visualize process gaps and get clear insights into operational inefficiencies or opportunities for improvement.

Explore the report in Accenture's thought leadership app, Foresight, and get a personalized feed of recent latest insights, data, case studies and more at http://www.accenture.com/foresight.

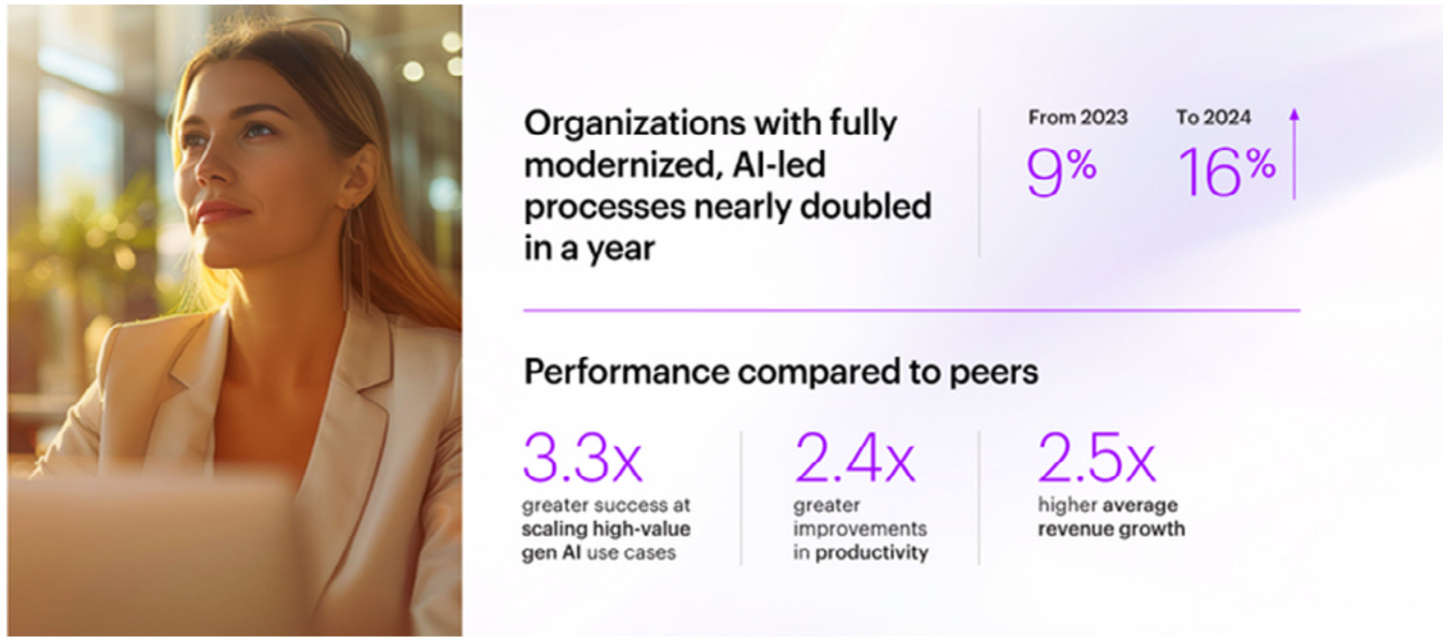

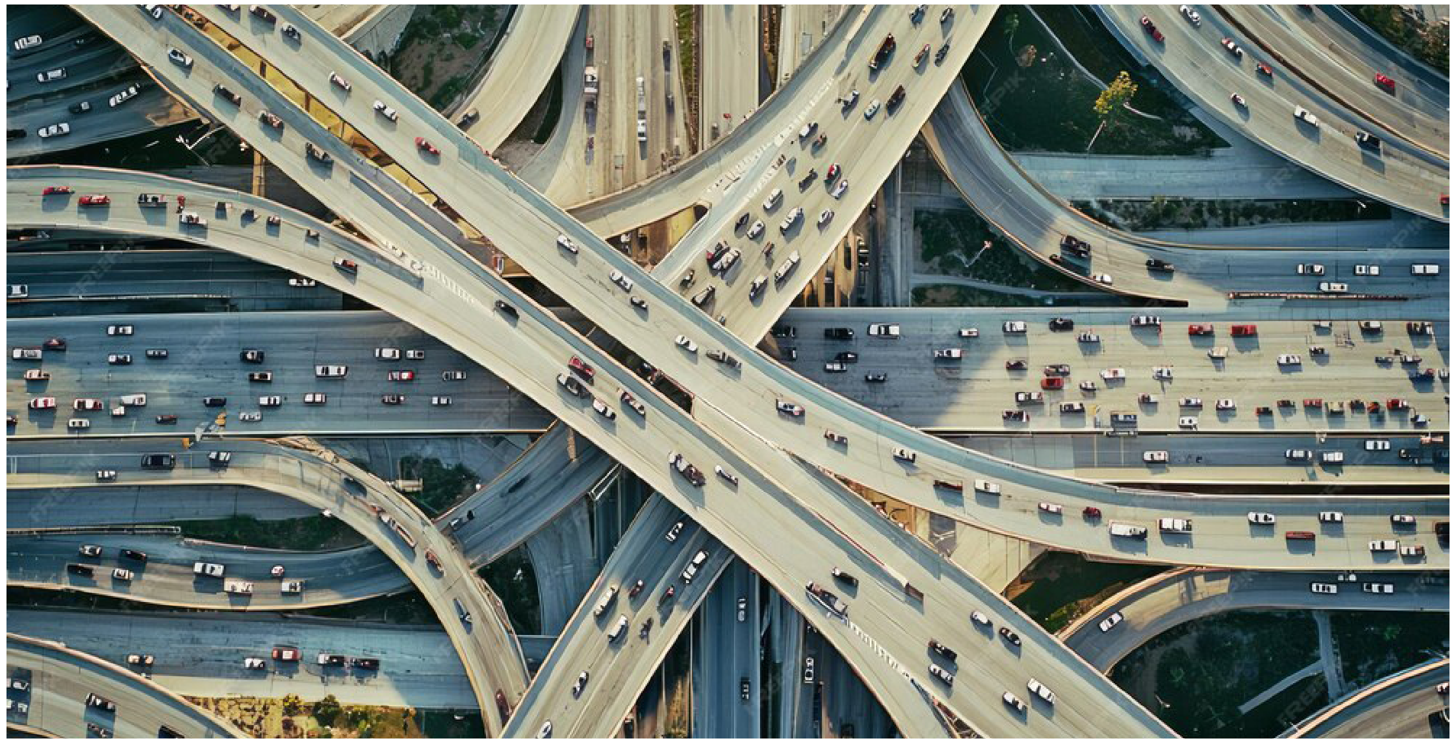

Impact on Data Governance with Generative AI

By Mark Restall

As artificial intelligence (AI) and machine learning (ML) technologies continue to transform industries and revolutionise the way we live and work, the importance of effective Data Governance cannot be overstated. With the emergence of generative AI, organisations are facing new challenges and opportunities in managing their data assets. This paper (over two parts) explores the intersection of Data Governance and generative AI, examining the traditional Data Governance model and its evolution to support generative AI.

What is Data Governance?

Data Governance is an applied framework that combines management, business, technical processes and technology to ensure that data is accurate, reliable, and secure. It involves tracking data throughout its lifecycle, from creation to disposal, to understand its meaning, control its use, and improve its quality. By building trust in data, Data Governance enables organisations to make informed decisions, comply with regulations, and maintain data security. This is achieved by setting internal standards, or data policies, that dictate how data is gathered, stored, accessed processed, and ultimately disposed of.

Business Benefits

- The biggest challenges organisations are facing to make themselves more "data-driven"
- Data is often not trusted.
- Hard to find and access data.
- Duplicated costs to the business as solutions developed in silos.
- Lack of traceability as to where data originated.
- Not having the right skills (data science/ architecture).

The greatest business benefits can accrue to an organisation when data is consistent, accessible and well-managed. Conversely, managing data effectively including understanding its quality, history, security, compliance and consent is important to reducing risks. These activities comprise data governance, and are critical to driving efficiency, productivity and trusted data, for better outcomes.

Purpose of Data Governance

The primary purpose of Data Governance is to achieve:

Shared Understanding: Provide a living framework for cross-organisational teams to have a common understanding of the data, who owns it, and how it should be handled.

High-Quality Data: Deliver high-quality data meeting metrics of high integrity, accuracy, completeness and consistency.

Data Profiling: Understanding of data based on factors such as accuracy, consistency/statistical of what it contains and timeliness.

Privacy and Compliance: Policies, standards and procedures drive technical and operational behaviours that ensure the systems meet the demands of government and industry regulations regarding sensitivity data and privacy, e.g. General Data Protection (GDPR), US Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standards (PCI DSS), Emerging AI regulations etc. Failure to comply has a significant impact on organisations.

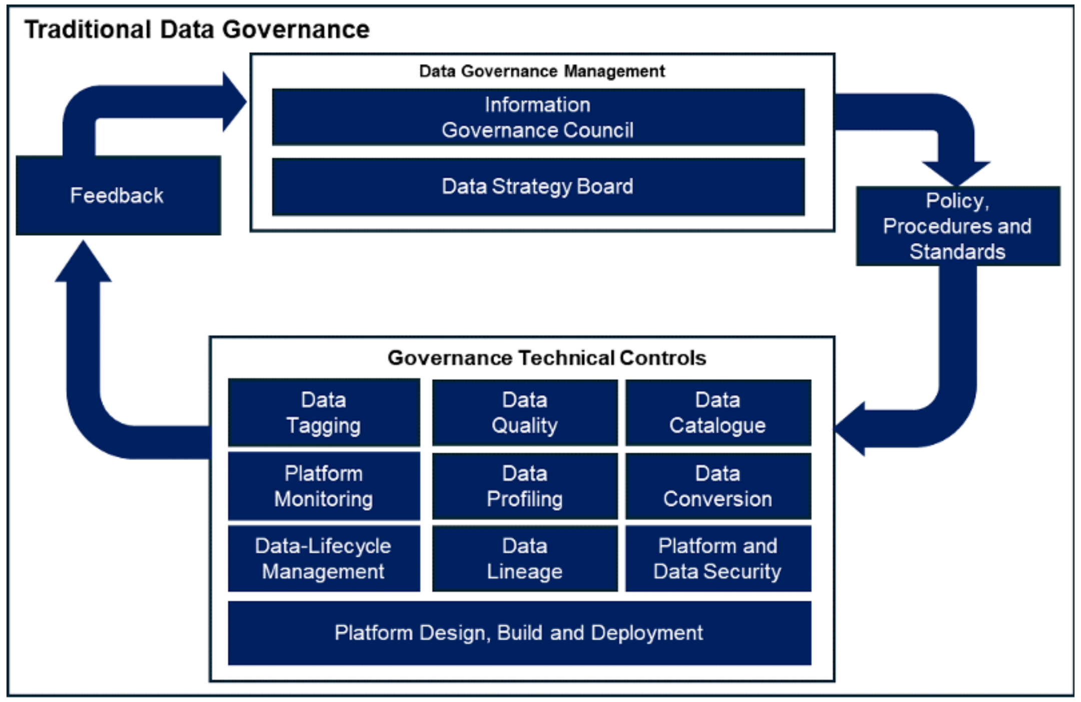
Traditional Data Governance
Data Governance Management
Information Governance Council
Data Strategy Board
Feedback
Policy, Procedures and Standards
Governance Technical Controls
Data Tagging
Data Quality
Data Catalogue
Platform Monitoring
Data Profiling
Data Conversion
Data-Lifecycle Management
Data Lineage
Platform and Data Security
Platform Design, Build and Deployment

While organisations approach Data Governance in different ways, there is a pattern which emerges, which we will call "Traditional Data Governance" illustrated above

Facilitate Feedback and Improvement: Provides a mechanism for human and technical feedback improve processes, policies, standards, technical controls which improves quality and security of the data.

Reduce Operational Cost: Process and storage of data is expensive, duplicate data sets and data workflows are not optimal, systems are not built using common standards etc. Data Governance has a role in addressing these areas and can help reduce the overall operational cost of the systems, support and underlying platforms.

Support Advanced Analytics and AI Use: Ensure high quality data to support advanced analytics, machine learning and generative AI initiatives. Analytics and Model trust driven, by trust in the data, drives adoption of models.

Monitor AI Use: Monitoring Machine Learning and generative AI is critical to detect results/out which cause reputational damage, incorrect behaviour, wrong advice, failure to comply or meet regulator enforced standards etc. Data Governance has a clear role in monitoring, reacting before and when things are going wrong.

Organisational Data Governance

It must be recognised that organisations differ significantly in operating model, purpose etc. As a result, the Data Governance models applied will vary significantly, and they may focus on specific elements, and be developing others, or not at all.

Data Governance Management

Data Governance, needs from an enforcement perspective, to be rooted in the organisation's senior leadership represented on the Information Governance Council (IGC), and through the day-to-day Data Strategy Board (DSB). Typically, a Senior Responsible Officer (SRO) such as a Chief Data Office or Chief Information Office is a role on the board of Directors and is ultimately responsible for data, how it's used, protected, worked on etc. within the organisation.

The IGC supported by DSB, is there to be the single data authority which owns, informs, monitors, enforces, creates, updates, retires etc. policies, procedures, standards and technical controls to support the business need.

Information Governance Council

The Information Governance Council is led/chaired by the SRO and sponsors, approves, and champions strategic information plans and policies. It owns the organisational mission for Data Governance within the organisation.

Data Strategy Board

This board handles the day-to-day issues regarding data and provides responses to those issues. The DSB owns and monitors the goals of information governance within the organisation, in line with those set by the IGC.

IGC and DSB Representatives

The IGC and DSB will have appropriate representatives to provide user, business, data, security and technical perspective etc. These representatives typically are:

Chief Architect: The Chief Architect can have oversight of the architectural, engineering and possibly support aspects of the organisation's platforms. Senior and more junior architects are represented in the Data Custodian persona. The Chief Architect is there to ensure common standards to architecture, component selection/ recommendation and best practice.

Data Owners: Data Owners are typically represented by senior officers / executives who have the authority to make decisions based on the information that belongs to them. They make decisions about the data to address the needs of their business function or the wider organisation. Data owners may not be involved in the day-to-day management of data. They will usually delegate data-related operational responsibilities to Data Stewards as appropriate; they cannot delegate their accountability.

Data Stewards: Stewards are managers of the data (information) and of its implications. Stewards are responsible for championing data, as well as maintaining and reconciling across different business units etc. the definitions for data, data quality, definitions and semantics, business rules and anything else delegated to them by the Data Owner.

(Continued over)

(From previous page)

Data Custodians: Data Custodians (typically technical teams) work closely with Data Owners, Data Stewards, and data security and protection teams to define data security and access procedures, administer access systems, manage the disposition of the data day-to-day (for example, management of a cloud data store), and provide backup and disaster recovery capabilities.

Platform and Data Security: Typically, a chief security officer would be represented on the IGC with senior security leads on the DSB. They would contribute to the overall Data Governance framework from a security perspective ensuring approaches and relevant standards are incorporated. At a DSB level they would represent the security teams working on a day-to-day level ensuring compliance and working to resolve any incidents.

Applied Data Governance

Policies, Standards and Procedures shape the entire data platform, data storage and data processing, from design to operation and through to decommissioning.

Below we cover the Data Governance technical controls (defined by the Policies, Standards and Procedures) which are used to improve the trust in the data.

These controls are not just for design and build, but throughout the support and the destruction of the data and platform.

Platform Design, Engineering, Deployment

Today, almost every component e.g. application, application code, infrastructure etc. of a platform can be deployed and configured through scripts. The collection of scripts, code and other artefacts are assets which can be stored and versioned.

DevOps tooling provides the ability to automatically and repeatably deploy, update and test the solution. The use of these standards and policies, in creation, reviews and revisions through feedback, it is possible to drive those assets to a state fully supporting the organisation's Data Governance goals for data and architecture.

Platform Security

Data security is embedded at design and throughout the platform lifetime, covering at-rest, on-the-wire encryption, strong role-based-access control (RBAC) model etc. Data Governance security standards, policies and procedures, for platform and data, enabling and controlling how architects, support, security etc. design, operate and monitor, ensuring the organisation is protected.

Data Governance Technical Controls

Data Governance technical controls embedded in the platform, monitor data, record and provide access to technical metadata about the data, monitor and improve data quality, track data processing etc. Platform Data Governance technical controls typically come in the form of:

- **Data Lineage:** Tracking/data flow and metrics of processing of the data.
- **Data Quality:** Checking the data against a defined set of rules.
- **Data Profiling:** Creating statistical metrics on the data itself.
- **Business-Glossary:** Stores a list of business terms and definitions used in describing data and its processing.
- **Data-Catalogue:** Stores list of fields, data type, descriptions, and other metadata describing is owner, location etc.
- **Quality-Correction:** Algorithms to correct errors in the data.
- **Data Conversion:** Basic algorithms to convert data formats or perform lightweight conversion of data values.
- **Data Tagging:** Added fields at ingestion to allow rollback, transformation, and deletion of selected datasets. Fields could be e.g. ingest-time, data owner, source, pipeline, security credentials etc.

It should be noted, that what is actually deployed will depend on the use-case (s) and organisational Data Governance requirements.

Data Life Cycle Management (DLM)

Data Lifecycle Management places data in a state for example data creation, data collection, data storage/stored, data processing, data sharing and usage etc. These states may change due to events, for example, requirement to destroy data, transfer to different systems, rights to use expiration, or simply change as part of navigation through the ingest data pipeline or processing operations.

The data state can be identified by its physical location in the platform layers, attributes in the Data Catalogue, specific Data Tagging etc. which requires different handling to move to the next state as defined by Data Governance policies and procedures.

Data Governance creates and enforces a DLM for data, ensuring the platform design, upgrade and destruction comply.

Platform Monitoring

Platform monitoring provides support teams with early warning of data and processing issues, predicts demand for optimal operation, capacity and demand management, and expense controls for example. Data monitoring using Data Governance platform technical controls, alerts the Data Custodians if there is increased quality-rule non-compliance from the source provider or perhaps problems in earlier stages in the platform, allowing investigation by Data Custodians and if required, feedback on standards, processes and procedures.

What Needs to Change in Traditional Data Governance?

Traditional Data Governance has served well over the years, however, there need to be increased capabilities or new elements added to support generative AI. Traditional Data Governance has:

Largely been focussed on structured data, however, to evolve to support generative AI, this needs to be expanded to support unstructured data derived from images, video, audio, and text for example at scale.

Limited/non-existent, support for model management, model history including generation, configuration, processes, data sources, model evaluation and test models to show a clear tree/lineage of the model, configuration and data at scale.

Only really tracked data preparation steps through a combination of DLM, data lineage and code version control. There needs to be a better understanding as to whether these preparation steps have introduced unhelpful artefacts into the model and creating a detailed versioned record of each step, including code, processing sequence and data used in these steps, in the model history.

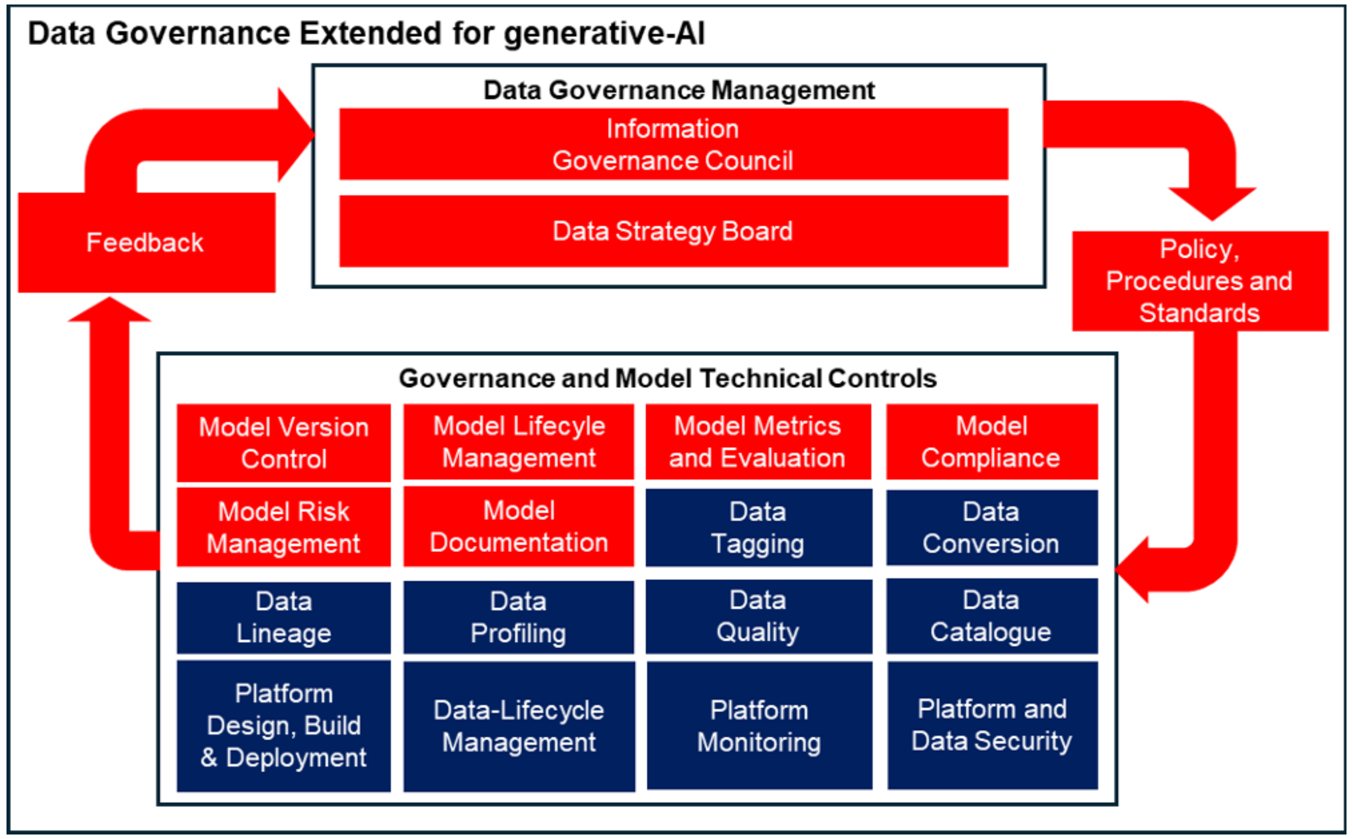
Data Governance Extended for generative-AI
Data Governance Management
Information Governance Council
Data Strategy Board
Feedback
Policy, Procedures and Standards
Governance and Model Technical Controls
Model Version Control
Model Lifecycle Management
Model Metrics and Evaluation
Model Compliance
Model Risk Management
Model Documentation
Data Tagging
Data Conversion
Data Lineage
Data Profiling
Data Quality
Data Catalogue
Platform Design, Build & Deployment
Data-Lifecycle Management
Platform Monitoring
Platform and Data Security

This diagram illustrates the evolution of the traditional Data Governance model to include the changes and added elements due to (applications for Machine Learning and) generative AI.

What is generative AI?

Generative AI is a type of artificial intelligence (AI) that can create original content, such as text, images, video, audio or software code in response to a user's prompt or request. Generative AI relies on machine learning models called deep learning models which simulate the learning and decision-making processes of the human brain. These models work by identifying and encoding the patterns and relationships in huge amounts of data, and then using that information to understand users' natural language requests or questions and respond with relevant new content.

Training a generative AI model

Training a generative AI model, begins with a deep learning model or foundation model, which is trained on huge volumes of data. This data is ingested, prepared, standardised, to create a neural network of parameters, calculations and data.

Traditional platforms are mainly focused on structured data. However, with generative AI the emphasis is towards multi-modal data. Therefore, the scope of Data Governance must have the policies, processes, and procedures to fully support multi-modal data (more biased towards unstructured e.g. text, images, video, audio,...), which adds different dimensions to how you undertake e.g. data quality checking, data profiling, data history/origin etc.

When driving new techniques of fact checking, for example, multiple source verification may be required. Furthermore, Data Governance can be used to improve explainability of the model outputs, capturing how a model was created (its data and steps to produce) and the influences that were used to shape its output.

This multi-modal data may ultimately live in different contexts, for example embedded/encoded within the model itself, but equally foundation models could use co-located or external data generated as models, to extend the base foundational model, shaping outputs for a particular organisation or give a new area of expertise.

Tuning

To improve accuracy, the model needs to be tuned to the specific task, typically using fine tuning or reinforcement learning with human feedback.

Fine tuning is where labelled model data specific to the types of questions or prompts that the end state model is expected to be asked, is used to train the model and refined to obtain the corresponding correct answers in the desired format.

Reinforcement learning is where human users respond to generated content with evaluations that the model can use to refine its response through re-training.

Data Governance should ensure that data sets, metadata, and human/system feedback responses correctly capture the model history change/lineage.

Monitoring and Refresh

Generative AI models (like all machine learning models) need continual system and user feedback monitoring – defining and evaluating performance metrics, to establish metric thresholds, or frequency when regeneration should occur.

External events such as a change in regulation, or a change in consumer behaviour, may invalidate some of the data sets used to build models forcing a refresh. The Data Governance maintained model and data lifecycle, may also be another trigger event for a refresh or some other action.

Data Governance may also provide guidance on the type of model that an organisation uses. Core foundation models are extremely expensive to produce (£10m-£100m), whereas:

(Continued over)

(From previous page)

Retrieval Augmentation Generation (RAG) models can be trained on smaller more focused data sets, extending foundation models, providing more accurate responses, are significantly cheaper, and are more easily tuned to ensure currency for the organisation, which can be fully traceable through Data Governance.

RedHat's InstructLab provides a very powerful way to enhance generative AI models with new content. InstructLab provides model upstream maintainers, with the required infrastructure resources, the ability to create regular builds of their open-source licensed models not by rebuilding and retraining the entire model, but by composing new skills into it, significantly saving time and cost, to ensure currency for the organisation – which can also be fully traceable through Data Governance.

What Should Change in Traditional Data Governance?

So, what should change? In summary, it's the transparent traceability of the elements that go into the process of creating and refreshing the model, combined with monitoring, to ensure that it is compliant, and returns accurate results without favour.

Data Governance Management

Data Governance Management (IGC and DSB) need to create policies, procedures and standards for data preparation, training, tuning, compliance, and lifecycle in the creation and update of models which take into account and monitor that the underlying data may have changed, the world that model exists may have changed, or other factors, requiring different responses than before.

Data Governance Management feedback mechanisms should be updated to ensure that the new technical controls to monitor models and human feedback are used effectively in shaping policies, procedures and standards to deliver the most accurate models, increase trust and reduces risk to the business.

Within an organisation, there also needs to be a Senior Responsible Owner (SRO) who is the designated Model Owner (potentially, distinct from the Chief Data/ Information Officer) and is responsible for all aspects of the model generation and running. The Model Owner will be a member of the Information Governance Board or Data Strategy board (as outlined in the earlier blog).

Model Owner

Generative AI models are different to machine learning models that have gone before, they use knowledge to create content, advise and sometimes make decisions. The Model Owner for key generative AI models will decide on what happens to the model, strive to deliver the best outcome, protecting the organisation from any legal ramifications and reputational damage.

The following types of Data Governance controls can be used, leveraging and extending some from the traditional framework and adding new ones:

Model Version Control

To build trust and ensure compliance and tracking of model code, data preparation code, test, and generation data sources used for model training should be fully captured under version control.

Model Life Cycle Management

Data Governance Model Lifecycle Management tracks the model during its lifetime and assigns it to a stage in the lifecycle e.g. Model Proposal, Model Design, Data Preparation, Model Building, Model Training and Tuning, Model Testing, Model Deployment, Model Deployment Testing, Model Re-Training and Tuning, Model Destruction etc.

Generative AI models differ from other digital assets in so far as their "knowledge" may have a time currency. A digital photo does not expire, the photo is a true and accurate reflection of the scene that was captured at that time. Its purpose is not to evolve how the scene evolves and nor can it. Generative AI models may require refresh to maintain its currency (due to events like regulatory change, failure to meet new metrics, user feedback, business feedback etc.), triggering a move to a different stage (e.g. Model Re-Training and Tuning, or Model Destruction) in the model life cycle.

Model Metrics and Evaluation

The Model should be evaluated against established model metrics agreed by the IGC and DSB, during testing and continuously during its execution. The key aims are to establish the model's performance, fairness, and stability. The following could be used to monitor generative AI:

Data Drift

Data Drift is where the statistical properties of the target variable (key input data) or input features change over time. By comparing (with the metrics) the model results with historical data, we can see if the results still reflect the expected historical data. If there is a difference, then this is due to data drift.

Model Performance

As part of model training, we established metrics for the model performance using a dataset or test dataset. Throughout its lifecycle, the model is evaluated against these metrics to ensure that it is working as expected.

Model Fairness

Model Fairness metrics ensure that models make predictions without introducing or perpetuating biases or discrimination.

Model Explainability

Model Explainability is the ability to interpret and explain how the model arrives at a given output.

Feedback

Feedback such as user feedback needs to be captured and incorporated which give evidence for change in the metrics above.

Model Compliance

Generative AI models are being used today to check for document regulatory compliance and for other tasks where validation against Government, Industry, or other body regulations are required.

Regulation can change quickly and many large organisations are affected by a wide range of regulations. Keeping track, and more importantly, deciding when to update is key to compliance.

Many of these tasks were done manually, but tooling is available today to support automated compliance checking to identify and enforce up and coming regulations.

Automation allows for recording of these changes in the Data Governance and triggering model changes ensuring that the organisation is up to date. Metrics would also need to be created so that when the model is not compliant it is flagged.

Model Risk Management

We are beginning to place huge levels of trust in our generative AI and ML models. As they become more sophisticated, they drive new risk profiles and complexities.

The Data Governance tooling using the approaches described can minimise privacy and copyright violations, as well as incorrect data which leads to false, misinterpreted, misleading or simply wrong outputs.

Data Governance tooling can also work with the test/ monitoring tools to set alerts to detect when specific metrics (human feedback, model fairness, bias, drift and specific model characteristics) are out of tolerance and then inform or re-train (refresh) the model to correct the issue.

Model Documentation

Increased transparency brings increased confidence in the model validation processes. It also supports and increases the explainability of AI for regulators, auditors and consumers.

It is not a great stretch to see the concept of a generative AI passport showing the family tree of data and processing being provided with each model.

Extended Traditional Data Governance Features

Data Tagging, Data Conversion, Data Lineage, Data Profiling, Data Quality and Data Catalogue etc. will be extended to support multi-modal data which means incorporating new metadata for data source types, results of more detailed profiling covering sentiment analysis, feature detection etc. with at scale data sets.

Data Conversion and Data Quality

Data Conversion and Data Quality needs to be more carefully handled to ensure any change in the content during data preparation and generation, does not adversely affect the information in the data and the model output. The processing code, data, data pipeline, and testing needs to be recorded as it tells the story of how the data is/has changed. These elements and flow would be recorded as part of the model version control, model metrics and evaluation, model documentation and data lineage.

Platform Monitoring

Existing platform monitoring will need to be updated to support continuous evaluation of the model using the Model Metrics including user and other relevant feedback.

Note: The scale of the training data (which in some cases is multi-Petabytes) means that practically (not from a technical perspective) tracing every data set from source could be difficult, time consuming and very expensive. However, with a world of litigation and copyright infringement, demands for increased trust in the models, detection of bias, reducing risk of reputational damage, the ability to support increasingly strenuous model validations etc. I see this transparency becoming a mandatory part of the generative AI toolkit.

Data and the generative AI Sceptic

One of the key roles of Data Governance is to communicate to the communities that use or consume the data. Traditional Data Governance would provide reports on, Data Quality, Data Profiling and Data Lineage showing the state of the data and where it has come from.

In the world of generative AI this becomes even more important as they are using the tools and creators, advisers and sometimes decision makers, the job of the Data Governance Board with the Model Owner is to be a sceptic of generative AI, advising clearly where and where it cannot be used for different use cases, with active feedback (human and technical) based on the current output, advising on improvements.

Generative AI As Part of The Solution

As well as presenting new opportunities for organisations to leverage their data in ways that previously were not economically possible, generative AI can also be part of the solution, as a crucial tool for helping automate a lot of the routine processes that would otherwise be entirely reliant on human input.

The main barrier to data governance at enterprise scale is effective user adoption due to time commitments and changes in organisation structure over time, so it makes sense wherever possible to explore the benefits of automation. Generative AI can support a number of key areas for Data Governance, specifically:

- Data Quality – automation of DQ checks, identification of errors and validation of data against pre-defined rules
- Data Cataloguing and Metadata Management – automation of the creation and management of data catalogues, making it easier to discover, access, and understand data assets and keep them up to date

Automation is crucial because it has the potential to reduce manual effort, freeing up resources for more strategic activities, improves quality and consistency enabling more informed decision making, and support organisations to meet regulatory requirements and standards – reducing the risk of non-compliance.

IBM has been supporting organisations to use generative AI to automate data governance processes that until recently could only be carried out by people. For example, reducing the time to collate and update metadata on image files not only significantly reduces time, but makes the exercise plausible and sustainable at an enterprise scale.

Conclusion

This article has hopefully helped describe the evolution of Data Governance in a world of generative AI. Today, we live in uncertain times, models have been created built on untrusted data sets, there is a lack of transparency in the generation and tuning, and as a result exhibit bias, lack accuracy etc.

Unfortunately, due to poor metrics, the true state is not really known, leading to issues for the consuming organisation which may well impact their consumers.

Tools are emerging and established in the market to address these challenges, tools like IBM watsonx.data Governance, which provides significant capabilities today addressing Model Risk, Model Compliance and Model Lifecycle supported by comprehensive Data Governance capabilities.

RedHat's InstructLab, further drives the cost and time down, to deliver the business freshness of generative AI models, and when combined with strong Data Governance can go along way to building trust in this technology.

In addition to the tooling, Data Governance frameworks need to evolve to address these challenges as outlined in this blog.

Is your business ready?

Originally published HERE

Mark Restall is Executive Architect, Data Technology and Transformation, IBM Consulting. (Many thanks to, Dr. Roushanak Rahmat, Hywel Evans, Joe Douglas, Dr. Nicole Mather and Russ Latham for their review feedback and contributions in this paper.)

Metadata driven information management in MS Document Management Systems

By Gerard Rooijakkers, Auckland Transport

Some of the biggest challenges in achieving metadata-driven information management are:

- developing **with** users the right type and number of metadata descriptors and deciding on standard terminology for the metadata values, basically developing the ontology and business taxonomy for the business. There are some standards and models available BUT the metadata must make sense for the business specific context.

- Helping users move away from (sub)folder structure is the challenge. Relying on auto-populating metadata tags when they upload documents can make the capture of information in SharePoint, MS Teams and OneDrive less onerous, it will make saving information easy, saving them time.

- Educating users in how to find information in any of MS document management applications and how to use filters when searching for information from the document management system is pivotal: Change Management is a must do investment. Demonstrating before and after auto-populating metadata tags with the information they save is an effective way to illustrate to the user of the effort and time saved by the implementation of such intelligent tool.

Artificial Intelligence & Machine Learning Technology to automate metadata tagging

The ever-increasing volume of information in the workplace comes from both the creation, receipt and aggregation of information in the organisation. Knowing what it is and where to find it is a crucial step in managing information. Users can easily access the information through a federated search, which can be enhanced with the use of filters to achieve better search results.

The benefit of automating the capture of information for users by using artificial intelligence (AI) is proven in many businesses that are applying these advanced tools. AI combined with machine learning is able to recognise content types and user patterns to add appropriate metadata tags which enables the delivery of the relevant information to users, Easy to Find.

Automating the retention and disposal process: life-cycle-management

Another vital tool to improve compliance and efficient management of information and data storage is a smart records management tool. The use of auto-classification to add metadata tags to information about their retention will assist and support an automated life-cycle-management.

The automation focusses on the creation of consignments of information for retention and disposal (RnD): archiving or purging. The human aspect in such an automated RnD process will remain in the electronic approval process, whereby the business is able to advise longer retention of the information than the legal requirement, the business purpose.

Conclusion

Implementing AI with Machine Learning tool to improve the capture of information created, received and aggregated will give the user that positive experience. This is the outcome you want to achieve as one of the Returns on Investment. Other outcomes are improved compliance and management of storage by implementing a smart records management tool.

By making the information findable and demonstrating what is possible, users will begin gathering confidence. Saving time and taking away frustration are big gains.

Arriving at such result requires the support from effective Change Management. The emphasis is on training staff and making them not only aware but also competent users of information. That is pivotal for an overall satisfactory success.

Easy to save, easy to find, easy to manage and complying with the NZ Public Records Act 2005 is the ultimate goal for every organisation in local and central government.

Gerard Rooijakkers is Corporate Information Manager at Auckland Transport. Originally published HERE

ABS explores Data Capture options for 2026 Census

The Australian Bureau of Statistics (ABS) is seeking information from organizations to potentially provide paper data capture services for the upcoming 2026 Census.

While most Australians have embraced digital methods for census participation, the ABS recognizes that a portion of the population still prefers or requires paper forms.

In 2021 over 78% of households complete their Census digitally. This equates to over 7.61 million online forms being submitted, and 2.87 million being submitted on Census Day alone.

In a recently issued Request for Information (RFI), the ABS aims to gauge market capabilities for end-to-end paper data capture solutions or specific components of the process.

The bureau is particularly interested in infrastructure, software, licensing, and support services that could be implemented at ABS premises in South East Melbourne.

This move comes after the successful processing of approximately 2.5 million paper forms during the 2021 Census. For the 2026 Census, the ABS is exploring various procurement models, including purchase and lease options.

The paper data capture operation for the 2026 Census is expected to handle a significant volume of forms, with the ABS estimating a need to process around 45,000 paper forms (1.1 million single sheets of A4) per day.

The process will require high-resolution scanning at 300dpi and advanced recognition methods such as OCR, OMR, and ICR to convert handwritten responses into electronic data. In addition there will need to be quality control measures to check poor image quality, completeness and validity of documents with ability to flag for rescan.

While the RFI is not part of a formal procurement process, it signals the ABS's proactive approach to planning for the next national census. The bureau emphasizes that this information-gathering exercise will help determine whether to proceed with one or more separate procurement processes in the future

RACGP wants Health Record overhaul

The Royal Australian College of GPs (RACGP) is calling on the Federal Government to overhaul My Health Record to make it fit-for-purpose for patients and GPs after a nationwide survey found 31% of GPs rarely or never use it.

The early findings from the RACGP's annual Health of the Nation survey come after a Productivity Commission report found My Health Record is 'plagued by incomplete records and poor useability,' with less than 2% of documents being seen by GPs.

RACGP President Dr Nicole Higgins said: "My Health Record can't fulfill its potential to be the one-stop store for Australians' health records without investment to improve its useability.

"It's a big job to improve our flagship national health data system, and we recognise that the Australian Government and the Australian Digital Health Agency are taking steps to this end. However, this work must continue to be prioritised because it will have significant benefits for Australians, and our health system.

"With improved useability, My Health Record will support better patient care, and better health for Australians.

"It will also make our health system more efficient and generate significant savings for the health budget – the Productivity Commission also estimated better use of electronic medical records systems can save up to $A5.4 billion each year by reducing the time patients spend in hospital, and $A355 million in duplicated tests in public hospitals.

"The potential savings in general practice and other health settings would also be substantial. This money could then be re-invested back into reducing out-of-pocket costs for patients to help in the current high cost of living climate.

Meta's cops Millions in GDPR Penalties

Meta, Facebook's parent company, has been fined €91 million ($150 million AUD) by the Irish Data Protection Commission (DPC) for storing user passwords as unencrypted 'plaintext' on its internal systems. The penalty comes five years after the company first acknowledged the security lapse.

The DPC, acting as the supervisory authority for the European Union's General Data Protection Regulation (GDPR), issued the fine along with a formal reprimand.

Deputy Commissioner Graham Doyle emphasized the severity of the breach, stating, "It is widely accepted that user passwords should not be stored in plaintext, considering the risks of abuse that arise from persons accessing such data. The passwords in question are particularly sensitive, as they would enable access to users' social media accounts."

The DPC's decision outlined four specific GDPR violations:

1. Failure to notify the DPC of a personal data breach concerning password storage (Article 33(1))
2. Failure to document personal data breaches related to password storage (Article 33(5))
3. Lack of appropriate technical or organizational measures to ensure password security (Article 5(1)(f))
4. Failure to implement measures ensuring ongoing password confidentiality (Article 32(1))

A Meta spokesperson responded to the decision, saying, "As part of a security review in 2019, we found that a subset of Facebook users' passwords were temporarily logged in a readable format within our internal data systems.

"We took immediate action to fix this error, and there is no evidence that these passwords were abused or accessed improperly. We proactively flagged this issue to our lead regulator, the Irish Data Protection Commission, and have engaged constructively with them throughout this inquiry."

Data capture solutions
that makes sense
What if information got where it needed to go... friction-free?

Australia has led the way regulating gene technology for over 20 years. Here's how it should apply that to AI

By Julia Powles, The University of Western Australia and Haris Yusoff, The University of Western Australia

Since 2019, the Australian Department for Industry, Science and Resources has been striving to make the nation a leader in "safe and responsible" artificial intelligence (AI). Key to this is a voluntary framework based on eight AI ethics principles, including "human-centred values", "fairness" and "transparency and explainability".

Every subsequent piece of national guidance on AI has spun off these eight principles, imploring business, government and schools to put them into practice. But these voluntary principles have no real hold on organisations that develop and deploy AI systems.

Last month, the Australian government started consulting on a proposal that struck a different tone. Acknowledging "voluntary compliance [...] is no longer enough", it spoke of "mandatory guardrails for AI in high-risk settings".

But the core idea of self-regulation remains stubbornly baked in. For example, it's up to AI developers to determine whether their AI system is high risk, by having regard to a set of risks that can only be described as endemic to large-scale AI systems.

If this high hurdle is met, what mandatory guardrails kick in? For the most part, companies simply need to demonstrate they have internal processes gesturing at the AI ethics principles. The proposal is most notable, then, for what it does *not* include. There is no oversight, no consequences, no refusal, no redress.

But there is a different, ready-to-hand model that Australia could adopt for AI. It comes from another critical technology in the national interest: gene technology.

A different model

Gene technology is what's behind genetically modified organisms. Like AI, it raises concerns for more than 60% of the population.

In Australia, it's regulated by the Office of the Gene Technology Regulator. The regulator was established in 2001 to meet the biotech boom in agriculture and health. Since then, it's become the exemplar of an expert-informed, highly transparent regulator focused on a specific technology with far-reaching consequences.

Three features have ensured the gene technology regulator's national and international success.

First, it's a single-mission body. It regulates dealings with genetically modified organisms:

> *"to protect the health and safety of people, and to protect the environment, by identifying risks posed by or as a result of gene technology."*

Second, it has a sophisticated decision-making structure. Thanks to it, the risk assessment of every application of gene technology in Australia is informed by sound expertise. It also insulates that assessment from political influence and corporate lobbying.

The regulator is informed by two integrated expert bodies: a Technical Advisory Committee and an Ethics and Community Consultative Committee. These bodies are complemented by Institutional Biosafety Committees supporting ongoing risk management at more than 200 research and commercial institutions accredited to use gene technology in Australia. This parallels best practice in food safety and drug safety.

Third, the regulator continuously integrates public input into its risk assessment process. It does so meaningfully and transparently. Every dealing with gene technology must be approved. Before a release into the wild, an exhaustive consultation process maximises review and oversight. This ensures a high threshold of public safety.

Regulating high-risk technologies

Together, these factors explain why Australia's gene technology regulator has been so successful. They also highlight what's missing in most emerging approaches to AI regulation.

The mandate of AI regulation typically involves an impossible compromise between protecting the public and supporting industry. As with gene regulation, it seeks to safeguard against risks. In the case of AI, those risks would be to health, the environment and human rights. But it also seeks to "maximise the opportunities that AI presents for our economy and society".

Second, currently proposed AI regulation outsources risk assessment and management to commercial AI providers. Instead, it should develop a national evidence base, informed by cross-disciplinary scientific, socio-technical and civil society expertise.

The argument goes that AI is "out of the bag", with potential applications too numerous and too mundane to regulate. Yet molecular biology methods are also well out of the bag. The gene tech regulator still maintains oversight of all uses of the technology, while continually working to categorise certain dealings as "exempt" or "low-risk" to facilitate research and development.

Third, the public has no meaningful opportunity to assent to dealings with AI. This is true regardless of whether it involves plundering the archives of our collective imaginations to build AI systems, or deploying them in ways that undercut dignity, autonomy and justice.

The lesson of more than two decades of gene regulation is that it doesn't stop innovation to regulate a promising new technology until it can demonstrate a history of non-damaging use to people and the environment. In fact, it saves it.

Julia Powles is Associate Professor of Law and Technology; Director, UWA Tech & Policy Lab, Law School, The University of Western Australia and Haris Yusoff, Research Associate at UWA Tech & Policy Lab, The University of Western Australia. This article is republished from The Conversation under a Creative Commons license. Read the original article.

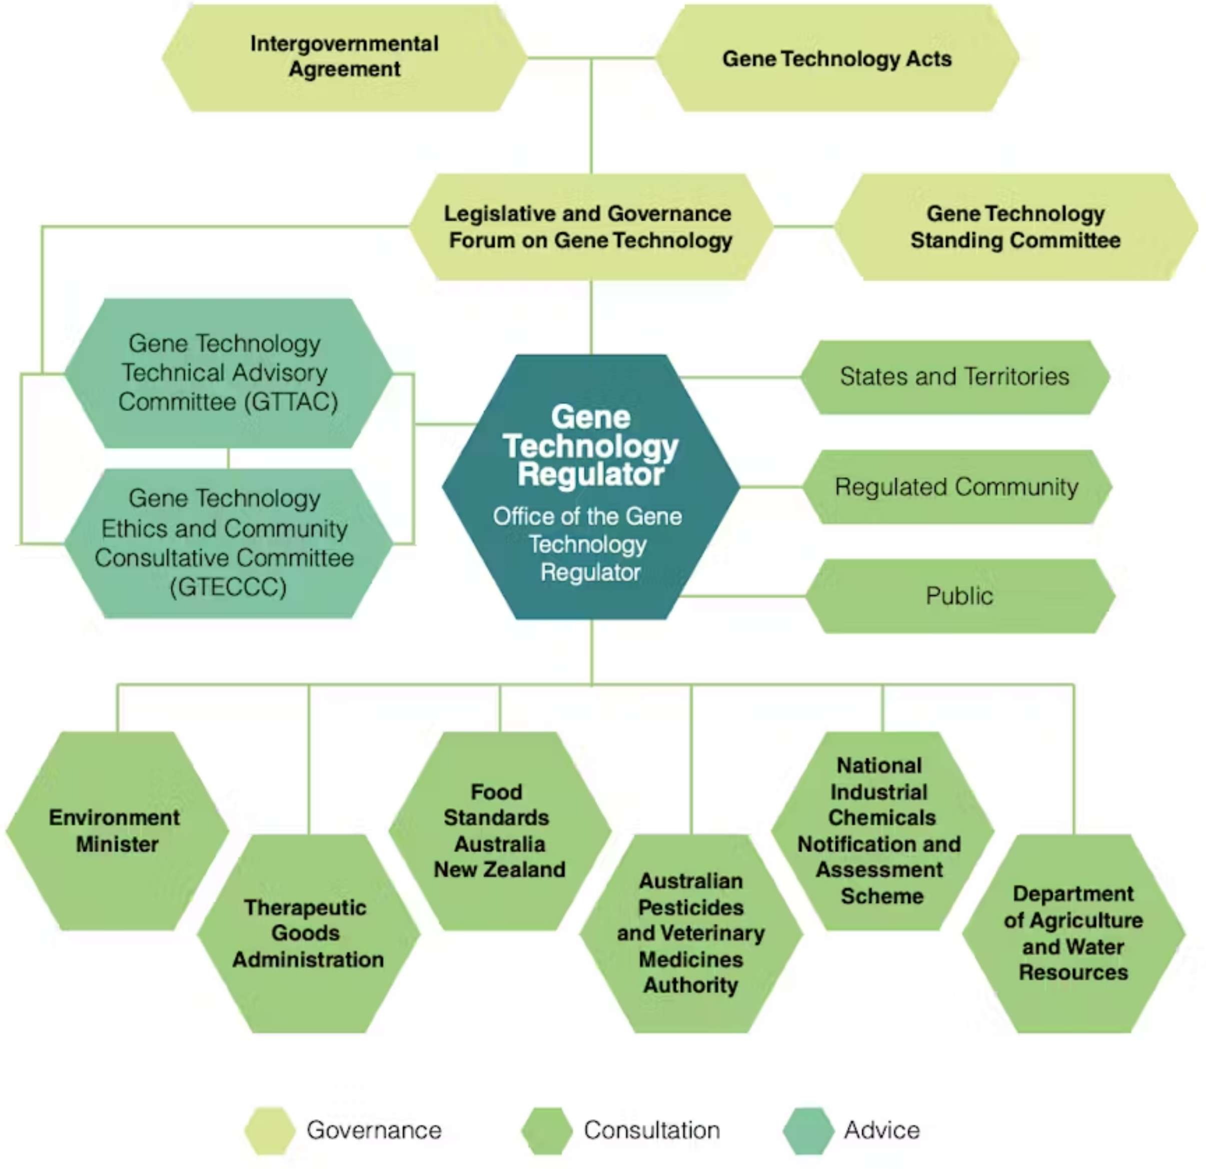
Intergovernmental Agreement
Gene Technology Acts
Legislative and Governance Forum on Gene Technology
Gene Technology Standing Committee
Gene Technology Technical Advisory Committee (GTTAC)
Gene Technology Ethics and Community Consultative Committee (GTECCC)
Gene Technology Regulator
Office of the Gene Technology Regulator
States and Territories
Regulated Community
Public
Environment Minister
Therapeutic Goods Administration
Food Standards Australia New Zealand
Australian Pesticides and Veterinary Medicines Authority
National Industrial Chemicals Notification and Assessment Scheme
Department of Agriculture and Water Resources
Governance
Consultation
Advice

The Gene Technology Regulator has a sophisticated decision-making structure. Office of The Gene Technology Regulator, CC BY

Automate with Confidence

By George Harpur

The goal of Intelligent Document Processing (IDP), or indeed any form of document automation, is no-touch or 'straight-through' processing, eliminating the need for time-consuming and often tedious manual intervention. Automation is easy, but *accurate automation* is not: historically, many IDP systems have failed to achieve this key outcome.

Accurate data is always important to some degree, and often the requirement is explicit in the form of a Service-Level Agreement (SLA) between a Document Processing Outsourcer and their client.

Say, for example, the SLA requires 99% accuracy in a classification task. If the automation alone is 95% accurate then on average 4 errors per 100 documents must be corrected to meet the SLA. Without knowing *where* the errors are, all 100 documents will have to be checked, and most of the efficiency gains from automation will be lost.

This, sadly, is how too many IDP systems are used today – often, document processors have reverted to 100% review after being bitten by inaccurate data and unhappy customers.

Confidence to the Rescue

The solution is for the IDP system to indicate where it is struggling to produce an answer, and therefore guide any manual checking to only a subset of the documents. In our above example, if the system can mark 10 classifications as being 'low confidence', and at least 4 of the 5 actual errors are within those 10, then we will be able to achieve the SLA while only needing to check 10% of the documents – 90% can now be fully automated.

Note that, for simplicity, we're assuming here that a human validating data is always accurate, and in reality that's not likely to be the case – in practice the IDP system will usually need to over-achieve the accuracy targets to offset human errors.

Can You Trust Your Confidence?

On the face of it then, it might seem that any system that can offer a confidence score is all we need, but unfortunately it's not as simple as that – a confidence value is only useful if it's a good predictor of where the actual errors lie, and many are not. To return to our simple example, if 2 or more of the errors occur in the high-confidence results (such errors are often called 'false positives'), then we're back to missing the SLA unless we manually check every document. We need to look more deeply at how confidence scores can be used, and how to assess their worth.

Confidence scores from a Machine Learning (ML) system are often expressed in the range 0 to 1, or as a percentage. You should be very wary though of assigning any particular meaning to a percentage unless you've had the opportunity to calibrate what it means in practice – a confidence of "99%" certainly sounds good, but it does *not* mean that the data is 99% likely to be accurate. Put another way, just because the system is confident doesn't mean it's right!

Confidence Thresholds and Accuracy

Although we ultimately care about accuracy (the percentage of results that are correct), not confidence, a reliable confidence score does give us a way to drive up accuracy while still retaining the benefits of automation. If we have a confidence score in the range 0-100, we can put a **threshold** on this score and manually review and data below the threshold. This value needs to be chosen carefully: too high and we will let too many errors through unseen; too low and we will need to review a large proportion of the documents and lose the automation benefits. This trade-off is best shown graphically. The graph below shows an example of classifying a test set of 1,000 documents from the insurance industry into 24 types.

The different points on the graph show different thresholds, with a high threshold on the left, decreasing to zero on the right. The leftmost point shows 100% accuracy (because there are no confident mistakes), but also 0% automation because *everything* has been marked for review (rejected)! By contrast, the rightmost point shows 100% automation (nothing is rejected) and 95% accuracy based on the 'best guess' result.

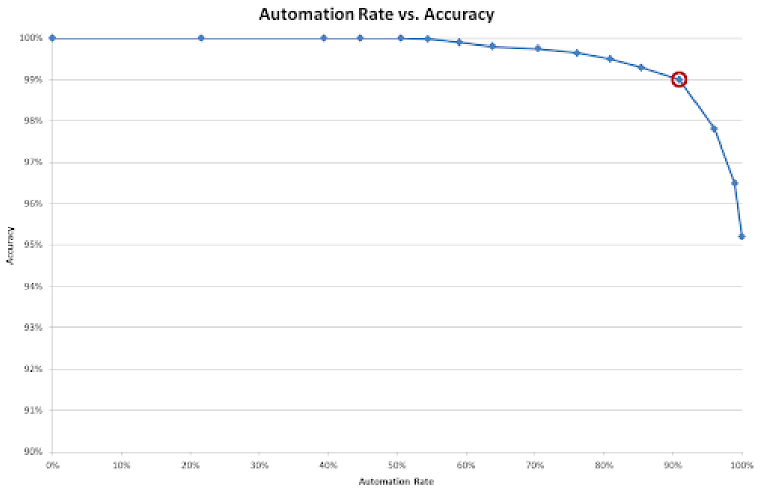

In between we see the benefits of a reliable confidence score – most of the incorrect classifications had low confidence, and in this instance manually reviewing only 10% of the documents would have enabled us to reach over 99% accuracy (by using the circled threshold value).

By contrast, let's look at a poor confidence score. In the graph below, I've deliberately sabotaged the confidence values by replacing them with random numbers. Although the endpoints are the same, the graph in between is very different, and now it would be impossible to achieve 99% accuracy without reviewing nearly all the documents.

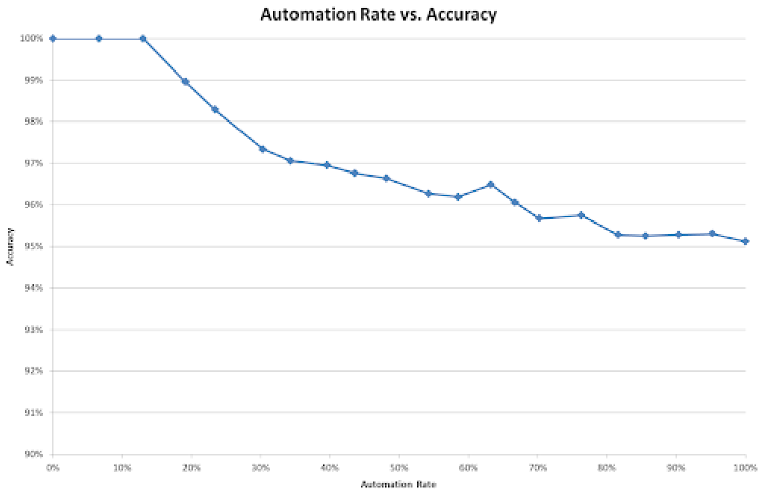

Although it's useful to understand the concepts behind confidence thresholds and their impact on accuracy and automation, most people do not have time or inclination to perform these experiments and set thresholds. A good classifier will 'self-tune' to a target accuracy by assessing its own performance and setting thresholds accordingly.

Track What You Hack

It's common to see vague claims about '90% automation' or '99% accuracy'. Firstly, armed with the information above, you now know that a single number is not the full story – an automation rate only makes complete sense if accompanied by an accuracy figure and vice versa. Secondly, it's important to ask 'percentage of what?' – a character-level accuracy will always be higher than a field-level accuracy, which in turn will be higher than a document-level accuracy.

Every project is different, so the only meaningful figures are those based on your own documents; all good IDP systems will provide easy ways to gather and track these. Here's an example dashboard from a project to extract hundreds of fields from traffic collision reports. The bars show the splits between field values that were captured correctly vs. those that required changes, and between those that were confident vs. unconfident, resulting in four categories overall. We can see a dip in performance on September 10th (less dark green, more red), which in this instance was because of a particularly challenging batch of documents that were poorly filled and scanned.

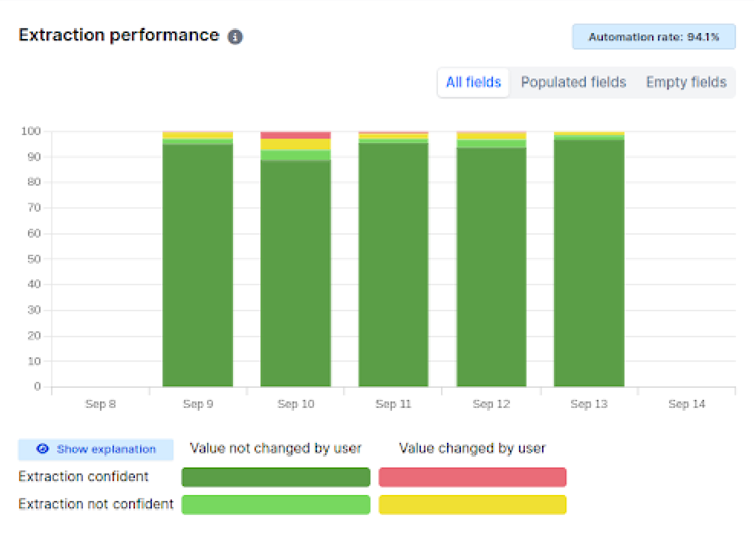

	Value not changed by user	Value changed by user
Extraction confident	🟩	🟥
Extraction not confident	🟩	🟨

A second view below shows a drill-down into the same figures for some of the individual fields on the documents. This provides actionable insights into where confidence levels may need to be changed or the configuration tweaked.

(Continued Over)

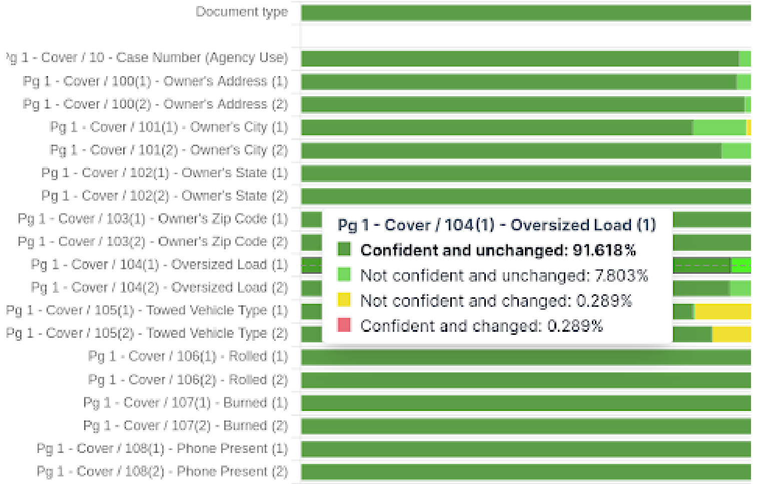
Document type
'g 1 - Cover / 10 - Case Number (Agency Use)
Pg 1 - Cover / 100(1) - Owner's Address (1)
Pg 1 - Cover / 100(2) - Owner's Address (2)
Pg 1 - Cover / 101(1) - Owner's City (1)
Pg 1 - Cover / 101(2) - Owner's City (2)
Pg 1 - Cover / 102(1) - Owner's State (1)
Pg 1 - Cover / 102(2) - Owner's State (2)
Pg 1 - Cover / 103(1) - Owner's Zip Code (1)
Pg 1 - Cover / 103(2) - Owner's Zip Code (2)
Pg 1 - Cover / 104(1) - Oversized Load (1)
Pg 1 - Cover / 104(2) - Oversized Load (2)
Pg 1 - Cover / 105(1) - Towed Vehicle Type (1)
Pg 1 - Cover / 105(2) - Towed Vehicle Type (2)
Pg 1 - Cover / 106(1) - Rolled (1)
Pg 1 - Cover / 106(2) - Rolled (2)
Pg 1 - Cover / 107(1) - Burned (1)
Pg 1 - Cover / 107(2) - Burned (2)
Pg 1 - Cover / 108(1) - Phone Present (1)
Pg 1 - Cover / 108(2) - Phone Present (2)

Pg 1 - Cover / 104(1) - Oversized Load (1)
- 🟩 **Confident and unchanged: 91.618%**
- 🟩 Not confident and unchanged: 7.803%
- 🟨 Not confident and changed: 0.289%
- 🟥 Confident and changed: 0.289%

Automate with Confidence

(From Previous Page)

It's important to remember that data not presented for review won't be checked for accuracy, which can distort the figures, so best practice is to perform full review on all documents for a brief period when first deploying a project in order to gather detailed statistics and build trust in the system. Thereafter, a small proportion (e.g. 1%) of documents can be sent to a quality control step for full review, which allows the performance of the automation, and the users, to be tracked over time.

The LLM Dilemma

Large Language Models (LLMs) are getting a lot of attention in the IDP world, and rightly so: on the one hand, they offer an almost magical way to pull complex data from documents with minimal configuration; on the other, they have a number of pitfalls if applied naively, and in particular around the issue of confidence and whether we can trust the data – are they just for toy projects and cool demos, or can they cut it in demanding high-accuracy environments?

At this point, most people have interacted with an LLM using ChatGPT or a similar interface. In this mode, it is easy to gloss over the limitations – if, for example, the LLM displays the well-known issue of *hallucination* (a.k.a. making stuff up), the user will often spot the issue and ask again or rephrase the question until it gives a better answer.

Chat, by its very nature, involves a 'human in the loop', and provides a natural nonsense-filter, but IDP is all about removing the human from the loop as far as possible. This is an example of the move from AI as *assistant* to autonomous *agent*. But how can we build sufficient trust in an LLM to use it in this mode?

Unfortunately, LLMs are notoriously bad at assigning meaningful confidence to their results. You might be tempted to simply ask an LLM 'how confident are you in that data?', but LLMs are trained to tell you what you want to hear: it will likely respond with a pleasingly plausible answer like "90%" or "95%", largely independent of the actual accuracy!

Those familiar with the technicalities of LLMs will be aware of 'log probabilities' which, with a suitable integration, give an indication of the confidence in each token (~word) at output. These can be used to build an overall picture of confidence, but are not usually good indicators on their own.

The reason is that the transformer architecture on which LLMs are built is based on a whole stack of layers, and while log probabilities give a reasonable assessment of confidence in the final word selection at output, the important 'decisions' about a particular classification or piece of data may have been made many steps previously and will be lost in the final probabilities.

Here's a simple analogy: say you're asked to classify a document, and as you form your final answer you're 50-50 on whether to call it 'letter' or 'correspondence' - this is equivalent to the probability you're likely to get from an LLM. However, the more important decision was one made several steps earlier in your thought process – is it correspondence at all, or actually a legal agreement written in the form of a letter? In an LLM, that kind of judgement-call is unlikely to surface in the output probabilities. A person would recognise that the initial decision was more important, and report their confidence accordingly, but that requires a level of self-reflection of which LLMs are not yet capable.

So what to do? At Aluma, we love using LLMs when they're the right tool for the job, but we don't use them exclusively, and we always use them with care. Here are some examples of the techniques we use to build a reliable measure of confidence in an LLM's output:

1. Cross-check the results against the original document text (hallucination filter);
2. Merge in the OCR/ICR confidence (poor-quality input gives poor-quality output);
3. Run the same LLM query multiple times (pick up random variations);
4. Have the LLM explain its working (hold it accountable);
5. Incorporate the LLMs' internal probability metrics ('white-box' integrations);
6. Run the query against multiple LLMs (two heads are better than one);
7. Cross-check the results using a different, e.g. rules-based, technique (two different heads are better than two similar heads);
8. Sanity-check the final results (validations help weed out lingering errors).

That's a lot, but the system automatically takes care of most of these behind the scenes, and our experience has been that with these methods in place, it *is* possible to use LLMs effectively in IDP, even in high-volume projects with demanding SLAs.

Summary (TL;DR Version)

Accuracy is always important in an IDP system, and confidence measures are key to achieving a high level of accuracy without a manual review of all the data. However, not all confidence values are trustworthy, so they should not be depended on without a proper analysis of how they relate to accuracy on a particular set of documents. It is also important to track accuracy over time to ensure that targets are still being met.

The latest generation of AI technology, based on LLMs, is powerful but has specific challenges in the area of accuracy and confidence. These models need to be enhanced with checks and balances to achieve the exacting requirements of a high-end IDP system.

George Harpur is Co-founder and CEO at Aluma. Originally published HERE

Data Breach Surge: A Wake-Up Call for Privacy Protection

The Office of the Australian Information Commissioner (OAIC) has reported a jump in reported data breaches in the first half of 2024 –at 527 the highest number in three and a half years.

Australian Privacy Commissioner Carly Kind said the high number of data breaches is evidence of the significant threats to Australians' privacy.

"Almost every day, my office is notified of data breaches where Australians are at likely risk of serious harm," she stated.

"This harm can range from an increase in scams and the risk of identity theft to emotional distress and even physical harm."

"Privacy and security measures are not keeping up with the threats facing Australians' personal information and addressing this must be a priority."

Malicious and criminal attacks accounted for over two-thirds of all breaches. The health sector and the Australian Government were the top targets, accounting for 19% and 12% of all breaches respectively.

The MediSecure data breach, which affected approximately 12.9 million Australians, stands out as a stark reminder of the massive scale these incidents can reach. This breach alone impacted nearly half of the country's population, making it the largest in the scheme's history.

The Australian Government has introduced the Privacy and Other Legislation Amendment Bill 2024. This legislation aims to bolster the OAIC's enforcement capabilities and clarify existing security obligations for organizations.

However, Commissioner Kind emphasizes that further reform is still needed to truly safeguard Australians' personal information.

"The Notifiable Data Breaches scheme is now mature, and we are moving into a new era in which our expectations of entities are higher," Commissioner Kind said.

"Our recent enforcement action, including against Medibank and Australian Clinical Labs, should send a strong message that keeping personal information secure and meeting the requirements of the scheme when a data breach occurs must be priorities for organisations."

The full OAIC report is available HERE

Australia makes Honour Roll of Top 11 Data Breaches of 2024

The MediSecure ransomware attack which exposed the personal and health data of nearly 13 million Australians, has earned a place in the top 11 data breaches of the first half of 2024, according to an analysis by Kiteworks.

Kiteworks, which delivers data privacy and compliance for sensitive content communications through its Private Content Network (PCN), used its Risk Exposure Index to analyze the top 11 data breaches of the first half of 2024.

The Index is a tool designed to evaluate and prioritize data breaches based on their severity and potential impact. Detailed in Kiteworks' "Top 11 Data Breaches in 1H 2024 Report," it goes beyond traditional metrics such as the number of records exposed or financial costs incurred.

Instead, it incorporates a range of factors to provide a more nuanced understanding of breach severity, including the type of data compromised, the extent of exposure, potential regulatory penalties, and long-term impact on brand reputation.

In conjunction with the report, Kiteworks made a Risk Exposure Calculator available on its website. This tool allows organizations to input their own data and calculate their risk exposure score, helping them better understand and mitigate potential cybersecurity risks.

"In today's complex cybersecurity landscape, organizations need a more sophisticated approach to assessing and prioritizing data breach risks," said Tim Freestone, Chief Strategy and Marketing Officer at Kiteworks.

"Our Risk Exposure Index offers a standardized framework for quantifying and comparing the risks associated with different data breaches, enabling organizations to allocate resources more effectively and enhance their overall security posture."

Key findings from the application of the Risk Exposure Index to the top 11 data breaches of 1H 2024 include:

The healthcare sector remains a prime target, with Change Healthcare's breach topping the list with a Risk Exposure Score of 9.46 out of 10.

Ransomware attacks continue to pose significant threats, as seen in the high-ranking breaches at Change Healthcare and Synnovis.

The sensitivity of exposed data plays a crucial role in determining risk, often outweighing the sheer volume of records compromised. The National Public Data breach tops the list here with a 9.46 out of 10 score.

Third-party and supply chain vulnerabilities remain a critical concern, as evidenced by the breaches affecting AT&T and Ticketmaster.

The report also highlights the diverse nature of cyber threats, from sophisticated ransomware attacks to inadvertent data sharing due to tracking codes. This diversity underscores the need for a multi-layered security approach that addresses a wide range of potential vulnerabilities.

"Our analysis reveals that the impact of a data breach extends far beyond immediate financial losses," added Patrick Spencer, VP of Corporate Marketing and Research at Kiteworks.

"The Risk Exposure Index takes into account factors such as regulatory compliance, data sensitivity, and potential for long-term reputational damage, providing a more holistic view of the true cost of a breach."

The Risk Exposure Calculator uses the same algorithm as the Risk Exposure Index, considering factors such as:

- Number of records exposed
- Estimated financial impact
- Ransomware involvement
- Data sensitivity
- Overall severity of the breach
- Number of regulations impacted

By providing this tool, Kiteworks aims to empower organizations to take a proactive approach to cybersecurity, identifying potential vulnerabilities before they can be exploited.

The report concludes with actionable recommendations for organizations to strengthen their cybersecurity posture, including:

- Adopting hardened security measures tailored to protect sensitive content communications
- Implementing advanced encryption techniques for data at rest, in transit, and in use
- Deploying next-generation Digital Rights Management (DRM) strategies
- Enhancing third-party risk management practices
- Focusing on data sensitivity and compliance through robust governance frameworks

The "Top 11 Data Breaches in 1H 2024 Report" and Risk Exposure Calculator are available at https://www.kiteworks.com/risk-exposure-index/.

LLMs and RAG ushering the Next Era of IDP

By Dr. He Zhang

Document automation has become a key tool for enterprises to improve efficiency and reduce costs. With the rise of generative AI in late 2022, Large Language Models (LLMs) like GPT, Gemini and Llama have demonstrated immense potential in document automation processing.

These models haven't simply transformed data processing methods; they've completely reshaped the way documents are processed, enhancing both efficiency and accuracy, particularly in Intelligent Document Processing (IDP).

Now a new approach is on the rise that blends real-world knowledge with LLM capabilities to enhance the ways back-office operations are automated – Retrieval Augmented Generation (RAG)

What is RAG and how does it work with LLMs?

Imagine you have a huge library of books, and you need to write a report on a specific topic. RAG is the super-smart helper who can do two things:

Retrieve: It quickly searches through all the books in the library and finds the most important information about your topic. It's like having a super-fast reader who can pick out just the right facts you need.

Generate: Using this information, it creates a report in its own words. It doesn't just copy from the books but understands the information and explains it in a way that makes sense for your use case.

So, RAG is like combining a super-fast library searcher with a brilliant writer. It helps create new information by first finding the right facts and then putting them together in a helpful way.

It feels like using an LLM because it generates texts on domain-specific topics – as a domain-trained LLM would. However, unlike standard LLMs, which are limited to the knowledge they were trained on, RAG goes a step further.

It utilizes retrieval mechanisms to access and incorporate outside knowledge, delivering more appropriate and precise information. This is especially helpful for tasks requiring extremely specific or up-to-date information.

RAG, in short, is what truly puts the "self-learning" capability on the LLM table.

Can I use an LLM without RAG and still achieve accuracy and efficiency?

On their journey to intelligent document automation, organizations combine the power of LLMs with domain-specific knowledge, thus deploying the so-called Customized Language Models for standalone use.

Whether they're named Specialized, Customized, Proprietary, or Private, they all refer to the same thing – a domain-specific language model, smaller than an LLM, that excels at handling private data sets within a specific topic.

Despite their edge of expertise, customized LLMs bring some challenges to organizations which slow down or block end-to-end automation**:**

■**Cost:** GPU costs for training language models can be quite high, and maintaining a certain number of continuously running GPUs to meet fast response demands further elevates costs.

■**Scalability and versatility:** Building and maintaining such models demands substantial initial development and ongoing maintenance, combined with very close collaboration across departments in the long term.

■**Transparency and explainability:** When errors arise or adjustments are needed, pinpointing the issue can be difficult for both customers and technical teams, because LLMs, and especially customized ones, only provide the action and not the reasoning behind it.

■So, the answer to the question above is yes, you can. In fact, many organizations do just that, but if they truly want to make the leap towards autonomous back-office operations, they need something extra...

RAG + LLMs: The power combo for document intelligence and automation

Enhancing LLMs with RAG makes it possible to overcome their built-in knowledge limits, produce more informed and contextually rich content, and pave the way for a new era where AI-generated text is not just more accurate, but also more nuanced and tailored to specific needs.

This spills over into many operational excellence gains for organisations:

■**Data Processing and Cost-Effectiveness:** The RAG + LLMs combination reduces dependence on expensive hardware by optimizing data organization and prompting processes.

The versatility of LLMs allows for faster adaptation to different customer needs, reducing the need for customized development and further minimizing cost and time investments.

■**Efficient Processing of Unstructured Documents:** Offering superior parsing capabilities, image and text data understanding, multimodal LLMs deliver more accurate information extraction and data classification, critical functionalities in the IDP domain.

Traditional processing methods struggle with unstructured documents like complex invoices. Multimodal LLMs, however, demonstrate their strengths in handling such tasks.

For example, in invoice processing, multimodal models can directly analyse image content and combine it with text information to effectively identify and interpret line items and table data without the need for cumbersome pre-processing steps.

■**Enhanced Transparency and Explainability:** By showcasing the entire decision-making process and reasoning chain, customers can better understand how the model functions and the basis for its decisions, thereby improving user experience and customer trust.

For instance, when an RAG model assigns an invoice to a specific general ledger account, it not only displays the outcome but also summarizes the reasoning behind the classification, such as pointing out the specific document content and historical data patterns that led to the decision.

> Despite their edge of expertise, customized LLMs bring some challenges to organizations which slow down or block end-to-end automation

This capability significantly enhances the transparency of the entire process, allowing customers to understand the logic and rationale behind each decision.

■**Provide a better continual learning experience:**

The ability to quickly provide feedback on the model allows customers to see improvements instantly, unlike conventional model training methods that require more time, effort, and cost.

As you can see the potential of combining LLMs and RAG extends beyond traditional document processing tasks. As technology advances, new application scenarios continue to emerge, such as automating complex business processes, enhancing customer interaction, and providing real-time business insights.

The bottom line? At Hypatos we are confident that the future of intelligent document processing is all about teaming up LLMs with RAG.

The former have already shaken things up in document processing, but the latter is what will take things to new heights if you are looking for spot-on accuracy, cost efficiency, handling complex and unstructured documents, and learning on the fly.

https://www.hypatos.ai/

Dr. He Zhang is CTO at Hypatos. Originally published HERE ezescan making digital work

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

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

ENCOMPAAS

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

encompaas.cloud | enquiries@encompaas.cloud | 1300 474 288

DOCUVAN IMAGE and DATA SOLUTIONS

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 | (03) 8535 3200

OPEX

OPEX® Corporation is the industry leader in document and mail automation, providing innovative, unique solutions that help streamline processes, and set the standard for operational efficiency. This includes seamless mail opening and sorting as well as document imaging (scanning), which increases throughput, maximises efficiency, saves time and money, and provides better output. Since 1975, the family-owned and operated company has served as a trusted partner to clients around the world, with more than 1,500 employees continuously reimagining automation technology that solves the most significant business challenges of today and in the future. OPEX provides advanced document and mail automation solutions across numerous industries, including service bureaus, law firms, banks, medical and health organisations, forms processing and archival agencies, and government institutions. OPEX is headquartered in Moorestown, NJ, with facilities in Pennsauken, NJ; Plano, TX; France; Germany; Switzerland; the United Kingdom; and Australia.

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

newgen

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. Furthermore, Newgen has a robust partner ecosystem, including global system integrators, consulting and advisory partners, value-added resellers, and technology partners.

newgensoft.com/home-anz/ | info@newgensoft.com | 02 80466880

iCognition

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

Hyland™

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

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

Kapish

Kapish is a member of the Citadel Group (ASX:CGL).Citadel solve complex problems and lower risk to our clients through our tailored advisory, implementation and managed services capabilities. With over 250 staff nationwide and an ability to 'reach back' and draw on the expertise of over 1,500 people, we are specialists at integrating knowhow, systems and people to provide information securely on an anywhere-anytime-any device basis. Servicing both large and small, public and private sector organisations across all industries, our team of highly qualified staff have global experience working with all versions of Micro Focus Content Manager (CM). It is this experience coupled with our extensive range of software solutions that enable our customers and their projects to be delivered faster, more cost-effectively and with more success. At Kapish we are passionate about all things Content Manager. As a Tier 1, Micro Focus Platinum Business Partner, we aim to provide our customers with the best software, services and support for all versions of the Electronic Document and Records Management System, Content Manager. Quite simply, our products for CM make record-keeping a breeze.

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

Kodak alaris

Kodak Alaris is a leading provider of information capture solutions that simplify business processes. We make it easy to transform documents and data into valuable business information and is where digital transformation starts. Kodak Alaris delivers intelligent document processing and information capture solutions that make sense. We exist to help the world make sense of information with smart, connected solutions powered by decades of image science innovation. Unlock the power of your information with our award-winning range of scanners, software and professional services available worldwide, and through our network of channel partners.

www.alarisworld.com/en-au | AskMe@kodakalaris.com| 1300 252 747

INFORMOTION SMARTER SAFER FASTER

INFORMOTION is an innovative professional services organisation specialising in the design and implementation of modern information management, collaboration and governance solutions – on-premises, in the cloud or hybrid. INFORMOTION's workflow tools, custom user interfaces and utilities seamlessly combine to deliver compliance, collaboration, capture and automation solutions that provide greater business value and security for all stakeholders. We can help you map and successfully execute your digital transformation strategy. Boasting the largest specialist IM&G consulting teams in Australia with experience that spans over twenty years, INFORMOTION consultants have a deep understanding of business and government processes and the regulatory frameworks that constrain major enterprises. Our compliance experience is second-to-none. INFORMOTION is a certified Micro Focus Platinum Partner and global Content Manager implementation leader. We are also an accredited Microsoft Enterprise Business Partner, Ephesoft Platinum Partner and EncompaaS Diamond Partner.

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

ABBYY fine tunes AI for document processing

ABBYY has previewed a suite of new capabilities for purpose-built AI at Ascend 2024, its inaugural technology and product showcase event.

A new multi-modal approach to zero-shot learning was previewed, known as ABBYY Phoenix, which leverages small language models and is purpose-built for document tasks.

Also announced was a new Secure Large Language Model (LLM) Gateway, along with enterprise-ready tools and applications that enable global enterprises to gain more value and insight into business-critical document processes.

The AI innovations are centralized in the new ABBYY Purpose-Built AI Center to help innovation leads and developers discover the array of tools available from ABBYY to accelerate application building.

The Purpose-Built AI Center serves as a centralized knowledge base for businesses to access information about ABBYY's suite of AI tools to develop next-generation applications. Here, enterprises and developers can learn about the range of capabilities offered by purpose-built AI from ABBYY, including how small language models are leveraged to deliver capabilities and accelerated time to value typically seen only by using LLMs.

By simplifying and accelerating the onboarding of new document types within enterprise business processes, organizations have more agility to adapt to evolving needs. This comes in addition to the 80+ document models already available out of the box.

By using ABBYY intelligent document processing (IDP) as a secure LLM gateway, businesses can command an LLM to extract data while simultaneously validating that the data is present in the document at hand. This empowers organizations to harness the power of general-purpose LLMs while significantly limiting hallucinations and increasing the reliability and trust of the output.

ABBYY also highlighted a collection of capabilities designed to improve efficient operations for large enterprises:

E-invoice processing: A streamlined, compliant processing of all invoice formats, ensuring adherence to diverse national standards.

Quality analytics: Tools for trend analysis of straight-through processing (STP) rates, continuously improving data extraction quality and performance.

Big data analysis: Enhanced scalability for process analysis, capable of handling complex, event-driven data with billions of data points.

"With the launch of our Purpose-Built AI Center, we are making it easier for our customer to find the right solution for their business needs and simplify their success in automation," said Maxime Vermeir, Senior Director of AI Strategy at ABBYY.

"ABBYY's innovations truly provide a response to the main concerns businesses have in the market today when leveraging AI and doing so at scale within an enterprise context."

ABBYY also announced its latest optical character recognition (OCR) performance improvements and tools equipping developers to more efficiently create cutting-edge applications. Highlights include:

Enhanced OCR capabilities with better memory consumption, improving document conversion quality and document structure detection.

New handwriting recognition (ICR) for multiple languages, document comparison modules, and updated processing profiles for data extraction, natural language processing (NLP), and archiving.

Simplified API support, as well as new and updated support for popular development languages such as .Net, Python, and more.

To learn more, visit the ABBYY Purpose-Built AI Center at www.abbyy.com/ai.

AI Navigator brings GenAI to the Desktop

Open-source data science platform Anaconda has released a desktop application that can run large language models (LLMs) on a local PC or Mac desktop. The free desktop application allows users of all levels to securely access and run over 200 pre-trained Generative AI models locally without sending any private information to external cloud services and infrastructure providers

As AI technologies continue to evolve rapidly, so have the challenges of securely adopting and managing these tools. Anaconda's AI Navigator, which is available for free download, addresses these challenges by offering access to a curated selection of over 200 AI models, tailored for a variety of tasks and device capabilities.

Users can access models to perform tasks such as document analysis, content creation, and data interpretation, tailored to their specific needs and computing resources. Agent creators can also leverage the local API servers to serve AI models locally for the creation and evaluation of AI agents.

Anaconda says early user interest has been especially strong for both code generation and debugging, with five of the top ten most downloaded models during the public beta coming from the CodeGemma and CodeLlama models.

Whether through the built-in chat interface or by integrating with external applications using the API inference server, AI Navigator allows businesses and individuals to experiment with LLMs in a way that suits their specific workflows.

In 2025, AI Navigator's centralized management tools will enable IT administrators to curate and govern the AI models their teams use, ensuring that only trusted, compliant models are deployed.

"AI Navigator's general availability signals our commitment to making AI accessible, secure, and powerful for all users," said Peter Wang, co-founder and Chief AI & Innovation Officer at Anaconda.

"Just as we democratized Python and data science to more than 45 million users, we're now empowering individuals and organizations to harness the full potential of enterprise-ready AI. As the foundation of Anaconda's OS for AI, AI Navigator ensures AI development is accessible, scalable, and secure for all."

"AI Navigator is just the beginning," said Steve Croce, VP of Product at Anaconda. "We're not only simplifying AI model access and management but laying the groundwork for a comprehensive ecosystem where users can build, govern, and deploy AI agents seamlessly.

"Our long-term vision is to make AI creation, distribution, and governance as accessible as any other everyday digital tool."

AI Navigator is now available for download.

GrooperAI simplifies AI usage

BIS has announced GrooperAI, the latest release of its intelligent automation software. It promises to make extracting and finding critical business data easier.

GrooperAI solves commonplace problems by simplifying and accelerating AI usage for information extraction, normalization, and consolidation.

Key Benefits of GrooperAI's Innovations:

Accelerated Decision Support: GrooperAI's no-code integration with leading LLMs generates human-level comprehension and AI-accelerated extraction, driving organizations to automate more complex tasks that require nuanced understanding and decision-making.

Improved Outcome: Easier setup, automatic prompting, and AI search capabilities significantly streamline workflows and reduce time to value.

Scalability and Flexibility: GrooperAI's adaptability allows businesses to scale automation efforts across departments and applications without extensive, time-consuming redeployment.

Future-Proofing: By incorporating AI-driven technologies, GrooperAI ensures that businesses can benefit from the latest advancements in automation without changing their production processes.

"These innovations mark a pivotal shift toward smarter, more adaptable, and financially transformative automation," said Rotelli.

"GrooperAI is at the forefront of the next generation of automation, where understanding and acting on complex data in real-time is key to staying competitive."

https://www.bisok.com/intelligent-document-processing/

Contracts to Payments simplified with SutiSoft

SutiSoft has announced the launch of a new platform to manage contracts. It allows users to create contracts, sign electronically, automatically generate invoices, and securely store documents - all within one interface.

"Switching between multiple systems to run your core operations adds layers of complexity, slowing down the process and increasing the chances of errors. We believe these critical workflows shouldn't be so cumbersome," saidN.D. Reddy, CEO of SutiSoft.

"That's why we envision a streamlined approach where every step - from contract creation to payment - can be managed effortlessly from a single interface. Our all-in-one platform transforms contract management from a fragmented, time-consuming process into a streamlined, efficient workflow.

"We're giving companies the tools they need to operate smarter, save time, and stay competitive in a digital-first world."

Key Features

Create Contracts Easily: Draft legally binding contracts using a user-friendly interface that supports customizable templates and workflows, ensuring consistency across agreements.

eSign Documents: Close deals faster with secure, legally compliant electronic signatures. The platform ensures seamless signing from multiple parties, with real-time notifications and progress tracking.

Generate Invoices Automatically: Once a contract is signed, the system automatically generates corresponding invoices, streamlining payment processes and ensuring precise billing.

Store and Manage Documents Efficiently: All documents—from contracts to signed agreements and invoices—are securely stored in a robust Document Management System (DMS), ensuring that only authorized people can access the documents with detailed audit trails.

https://www.sutisoft.com/suticlm/

Global Management Toolkit for GenAI

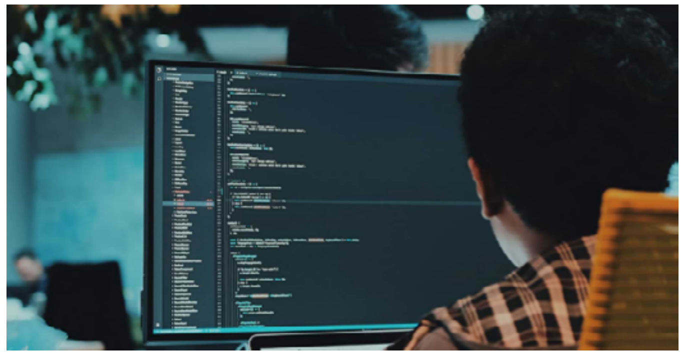

boost.ai, a developer of of conversational AI (CAI) solutions, has launched a new suite of tools designed to enhance the global management of guardrails and knowledge. The new toolkit, named Generative Action Management, will make it easier for enterprise customers to scale the use of Generative AI.

Building on the April launch of Generative Action, boost.ai's latest release significantly enhances the transparency and control of its GenAI-powered Virtual Agents (GenAI agents).

Without the proper guardrails and ability to manage knowledge in a centralized location, AI is a risk to the enterprise. For businesses operating at scale, or on track to achieve it, a lack of centralized oversight can lead to hallucinations, non-compliance, and potential misuse all compounded by the sheer volume of knowledge necessary to be effective.

boost.ai has directly addressed these issues by centralizing the two core components that enable the use of GenAI Agents in even highly regulated environments, guardrails and knowledge management. These new features empower enterprises to confidently deploy Generative AI at scale, effectively minimizing risk and ensuring that AI-driven customer interactions are accurate, compliant, and reliable.

"Just like human customer service agents, AI agents must be equipped with the right information to provide accurate and reliable responses. That's why it's crucial for AI managers to have the capability to manage and optimize the knowledge sources used by their AI agents," said Jerry Haywood, CEO of boost.ai.

"Before today, enterprises have had to work backward, decoding how LLM-powered AI arrived at a response once it was already delivered to a customer. Our toolkit arms enterprises with the resources necessary to utilize their knowledge as it is, and if necessary, reformat content in a way that is optimal for AI agents powered by generative AI. It provides a level of transparency into enterprise-ready GenAI that has yet to be seen."

With Generative Action Management, enterprises will gain access to the following new capabilities:

■ **Knowledge Management:** By gaining a clear and comprehensive overview of all knowledge being utilized by GenAI agents, users can easily confirm the correct knowledge is being tapped for each Generative Action taken by GenAI agents. Rather than tracking down issues on an individual topic basis, AI managers can validate knowledge relevancy and optimize formatting for the LLM, keeping GenAI agents up to date and reducing the risk of inaccuracies.

■ **Guardrail Management:** Enterprises at scale need GenAI agents that can accurately respond to potentially tens of thousands of topics and actions. Guardrail Management provides a centralized view of all guardrails in place, allowing users to control where and how guardrails are applied, identify which are mandatory, and enforce compliance from a centralized location.

■ **New Knowledge Types:** Knowledge Connectors allow knowledge from external systems, such as SharePoint and Salesforce to be connected directly, eliminating the need for file uploads or updates. By syncing documents from external sources, knowledge stays up-to-date and available for GenAI agents to utilize. In addition, Boost Documents provides users the capability to create and edit native documents directly in the platform, removing the tedious process of managing and re-uploading external versions of documents.

"Safe and responsible AI can only be achieved with strong guardrails that support today's needs, but are designed with the future in mind. With this launch, we've delivered enterprises a comprehensive toolkit to manage every guardrail and apply those necessary to ensure AI agents are safe and compliant," said Rasmus Hauch, CTO of boost.ai.

"Centralized management of guardrails is essential for ensuring complete oversight and control, empowering enterprises to enforce best practices, reduce risks, and confidently scale AI across their operations."

To learn more, click here.

Box in Alphamoon Acquisition

Box, Inc. has announced its acquisition of Alphamoon's AI-powered intelligent document processing (IDP) technology. Coupled with Box's recent acquisition of Crooze, this significantly expands the company's Intelligent Content Management (ICM) platform capabilities.

The integration of Alphamoon's technology with Box AI aims to revolutionize how enterprises extract value from their content. By combining advanced OCR technology with large language models like GPT-4 and Gemini, Box plans to offer enhanced metadata extraction and document structuring capabilities.

Aaron Levie, Co-Founder and CEO of Box, emphasized the significance of this acquisition, stating, "We are entering a new era of Intelligent Content Management. This acquisition represents a pivotal moment in our mission to revolutionize how enterprises derive value from their content."

The enhanced platform is expected to address various use cases, including automating metadata extraction from business documents, streamlining workflows based on extracted data, and integrating information with external applications like Salesforce.

"IDP remains a critical origination point for increasingly complex document workflow using intelligent capture paired with automation," said Amy Machado, Senior Research Manager, Enterprise Management & Workplace Strategies, IDC.

"To enable automation, organizations need a high level of accuracy, usability, and flexibility to support their diversity of documents and custom business requirements. With the Alphamoon acquisition, Box now has a complete end-to-end solution, unlike the pure-play IDP vendors."

Alphamoon's technology will provide Box with capabilities such as:

- Automating metadata extraction from business documents allowing users to leverage information stored within invoices, purchase orders, financial statements, contracts, leases, and more;

- Streamlining workflows based on extracted metadata, such as using metadata to understand if a legal contract is missing information, and alerting a legal assistant to complete it;

- Integrating extracted information with external applications like Salesforce. Today, the Box for Salesforce integration is one of the most widely used at Box and by leveraging enhanced metadata extraction in Box, enterprises can optimize core workflows, like loan processing;

- Enhancing industry-specific processes, such as analyzing clinical studies in life sciences or financial documents in the financial services sector.

Box has not yet announced availability and pricing for the new intelligent document processing capabilities.

Cloudflare Protects from AI Bots

Cloudflare has announced AI Audit, a new set of tools to help Web sites of any size analyse and control how their content is used by artificial intelligence (AI) models.

The company says it allows Web site and content creators to be able to quickly and easily understand how AI model providers are using their content, and then take control of whether and how the models are able to access it.

Additionally, Cloudflare is developing a new feature where content creators can reliably set a fair price for their content that is used by AI companies for model training and retrieval augmented generation (RAG).

Web site owners, whether for-profit companies, media and news publications, or small personal sites, may be surprised to learn AI bots of all types are scanning their content thousands of times every day without the content creator knowing or being compensated, causing significant destruction of value for businesses large and small.

Even when Web site owners are aware of how AI bots are using their content, they lack a sophisticated way to determine what scanning to allow and a simple way to take action. For society to continue to benefit from the depth and diversity of content on the Internet, content creators need the tools to take back control.

"AI will dramatically change content online, and we must all decide together what its future will look like," said Matthew Prince, co-founder and CEO, Cloudflare. "Content creators and website owners of all sizes deserve to own and have control over their content. If they don't, the quality of online information will deteriorate or be locked exclusively behind paywalls.

"With Cloudflare's scale and global infrastructure, we believe we can provide the tools and set the standards to give websites, publishers, and content creators control and fair compensation for their contribution to the Internet, while still enabling AI model providers to innovate."

Promised features include:

■ **Automatically control AI bots, for free:** AI is a quickly evolving space, and many website owners need time to understand and analyse how AI bots are affecting their traffic or business. Many small sites don't have the skills or bandwidth to manually block AI bots. The ability to block all AI bots in one click puts content creators back in control.

■ **Tap into analytics to see how AI bots access their content:** Every site using Cloudflare now has access to analytics to understand why, when, and how often AI models access their website. Website owners can now make a distinction between bots – for example, text generative bots that still credit the source of the data they use when generating a response, versus bots that scrape data with no attribution or credit.

■ **Better protect their rights when negotiating with model providers:** An increasing number of sites are signing agreements directly with model providers to license the training and retrieval of content in exchange for payment. Cloudflare's AI Audit tab will provide advanced analytics to understand metrics that are commonly used in these negotiations, like the rate of crawling for certain sections or the entire page. Cloudflare will also model terms of use that every content creator can add to their sites to legally protect their rights.

Existing Cloudflare customers can access the AI Tab from their dashboard today to review analytics for their sites and start controlling bots now. Site owners can visit https://www.cloudflare.com/lp/ai-value-tool-waitlist/ to join a waitlist to participate in the beta for price setting capabilities.

More information in this blog post at Start auditing and controlling the AI models accessing your content

Esker Transforms Order Processing

Esker has launched its newest Esker Synergy AI product update, Synergy Transformer. Designed to streamline order processing by optimizing data extraction, Esker's Synergy Transformer offers enhanced speed and accuracy using a custom-trained language model.

The model harnesses advanced Transformer technology, while training data is specifically tailored to the nuances of order language, ensuring precise and efficient data extraction.

Purpose-built for this specific use case, the Synergy Transformer is claimed to be faster at extracting large quantities of data from orders than large language models like ChatGPT 4 that focus on broader capabilities.

Integrated into Esker's solutions, the Synergy Transformer enables organizations to equip their CSR teams with AI capabilities, eliminating the need for extensive investments of time and resources in sourcing, evaluating, testing and refining alternative models.

The company says the updated Transformer provides a 6% improvement compared to the previous iteration, bringing the recognition rate to over 92%. Additionally, Synergy's smaller size is designed to be more sustainable and resource-efficient.

"I'm really excited that Synergy Transformer AI is now available to customers," said Aurélien Coq, Product Manager at Esker. "This new product feature further liberates CSRs by automating error-prone order data entry, freeing them to focus on strategic priorities."

https://www.esker.com/solutions/technologies/esker-synergy-ai

EXL launches Insurance LLM

A Large Language Model (LLM) built specifically for the insurance industry outperforms leading pre-trained models on accuracy across wide range of claims and underwriting related tasks, according to developer EXL.

According to Gartner, more than 50% of the GenAI models that enterprises use will be specific to either an industry or business function by 2027 - up from approximately 1% in 2023.

In internal studies, the EXL Insurance LLM achieved a 30% improvement in accuracy on insurance tasks, surpassing top pre-trained models, such as OpenAI GPT4, Claude and Gemini.

It was built by EXL AI Labs, using the full-stack NVIDIA AI platform, to support critical claims and underwriting-related tasks, such as claims reconciliation, data extraction and interpretation, question-answering, anomaly detection and chronology summarization.

EXL Insurance says the LLM was developed to address the highly specialized needs of the insurance industry, which has struggled to leverage off-the-shelf, general LLMs that lack fine-tuning of private insurance data and domain-specific understanding of business process operations.

Generic LLMs also fail to address the nuanced challenges faced by insurance companies during claim adjudication, leading to inefficiencies, high indemnity costs, claims leakage, longer settlement timelines, and increased compliance risks. By focusing exclusively on insurance-related tasks, EXL has incorporated its deep knowledge of the insurance industry and highly tailored proprietary data.

This level of specialization has become critical for ensuring accuracy, reducing cost and improving consistency in industry-specific AI applications.

"With 25 years of expertise in processing medical records data for bodily injury, workers' compensation, and general liability claims, EXL has developed curated data sets with domain-specific tagging, labelling, and question and answer pair creation for claims adjudication to fine-tune our models," said Anand "Andy" Logani, EXL's executive vice president and chief digital officer.

Specific tasks supported by the EXL Insurance LLM include the following:

■ **Structured and Unstructured Data Ingestion**: EXL Insurance LLM is able to aggregate and reconcile hundreds of thousands of de-identified medical records, claims histories, hand-written notes, call logs, and other claims and underwriting-related information.

■ **Contextual Classification and Triaging:** Data and insights extracted using the LLM are automatically categorized and fed into a wide range of core functions, ranging from claims adjudication to provider engagement to payment integrity to customer service functions.

■ **Conversations and Insights from Data**: Insights, question-answering and summary data drawn from the LLM empower faster, more accurate negotiations with providers, more robust assessment of anomalies and inaccurate payments and more personalized, real-time conversations with customers.

For more information about the EXL LLM for Insurance, visit here.

www.exlservice.com

Testing Framework for GenAI Models

Fisent Technologies has unveiled a new framework designed to help objectively evaluate the performance of various GenAI models for specific business process automation use cases. Fisent's *GenAI Efficacy Framework (GEF)* enables enterprises to measure, compare, and select the most effective GenAI models based on key metrics like accuracy, speed, cost, and consistency.

Fisent's GEF empowers enterprises to optimize their GenAI implementations by identifying the best model and parameters for each of their process automation efforts. This data-driven approach helps teams justify their GenAI choices and demonstrate a clear return on investment.

By evaluating models against real-world requirements, GEF mitigates the risk of suboptimal performance and ensures that organizations stay ahead of the rapidly evolving GenAI landscape through continuous assessment and optimization.

"The idea for GEF sparked as typical AI model evaluation methods, like those that measure Massive Multitask Language Understanding (MMLU), failed to balance the nuanced requirements of our customers' real-world automation decisions," explains Adrian Murray, Founder and CEO of Fisent.

"GEF offers a more pragmatic approach by evaluating the most important factors to any given application decision: accuracy, speed, cost, and consistency.

"Understanding these metrics allows enterprises to make more informed decisions about which LLM to employ for each of their process automation challenges."

Fisent's GEF includes a configurator that customers can use to evaluate the tradeoffs inherent in comparing their LLM options. The GEF configurator intelligently scores given requirements against available LLMs to produce a ranked list of the models expected to perform best for a specific situation along with numerous visual comparison charts for speedy analysis.

For example, using the configurator to increase the requirement for accuracy will adjust the rank-order of LLMs under consideration. What's more, other variables evaluated by the configurator adjust accordingly.

If there is a need for both speed and high accuracy, the cost variable associated with the best-fit LLMs is likely to increase. GEF is useful to Fisent customers when they initially implement Fisent BizAI against a specific process automation and again when evaluating new models or model upgrades.

Key benefits of Fisent's GEF include:

■ Comprehensive evaluation: Assess GenAI models across multiple variables, such as accuracy, speed, cost, and consistency.

■ Data-driven insights: Provide actionable recommendations based on objective metrics and statistical analysis.

■ Continuous optimization: Enable enterprises to monitor and improve model performance over time.

■ Ease of use: Streamline the process of evaluating and selecting GenAI models, even for non-technical users.

http://www.fisent.com

Hyland™

Hyland Content Innovation Cloud

Hyland has announced a new cloud platform to unify its content solutions — OnBase, Alfresco, Nuxeo and Perceptive Content. The Content Innovation Cloud is described as a unified content, process and application intelligence platform.

It will offer the potential to "AI-enable" enterprise content wherever it resides - including third-party repositories - resulting in more informed decision-making and strategic action, while ensuring governed and secured access to content and application of AI.

It will expand the functional capabilities of Hyland's Content Innovation Cloud - which already includes cloud-native services Hyland Insight, Automate and Credentials - with forthcoming advanced content services, intelligent document processing, knowledge discovery and intelligent automation

Hyland says it will deliver an intelligence layer that AI-enables content without the need to decommission or migrate systems.

"The Content Innovation Cloud will unlock the critical intelligence within our customers' vast array of content, transforming unstructured data into valuable insights that spark innovation and deliver unprecedented business outcomes," said Jitesh S. Ghai, Hyland president and CEO.

"Content is the foundation of digital transformation, and we're redefining the possibilities of both by transforming content management into content innovation."

Key components of the platform roadmap include:

■ **Content intelligence** that leverages advanced AI to analyze and interpret enterprise content wherever it resides, transforming unstructured data into actionable insights

■ **Process intelligence** that harnesses AI and advanced analytics to continuously monitor and optimize content-driven business processes, uncovering inefficiencies, automating tasks and enhancing operational performance

■ **Application intelligence** that uses AI to integrate content insights directly into business applications like Salesforce and Workday, enhancing user experience, expanding functionality and enabling data-driven decisions across industry-specific apps

To learn more, visit Hyland.com.

Dropbox Dash adds Universal Search

Dropbox has announced Dropbox Dash for Business - the latest iteration of its AI-powered universal search product - is now available in the US, with availability in international markets in early 2025.

It combines universal search, organization and sharing capabilities, and advanced content access control.

Modern work is more distributed and virtual than ever, with information scattered and siloed across a sea of browser tabs, cloud apps, and AI tools. As a result, knowledge workers spend too much time painstakingly searching for content across email, team drives, and cloud apps.

Dash addresses the constant friction teams feel at work by connecting with work apps to create a central hub to find anything in one place. And Dash comes with powerful content access and permission controls to guarantee company content is seen only by the right people.

Dash integrates with essential tools such as Google Drive, OneDrive, Notion, Asana, and more. It uses machine intelligence to improve search results and provide realtime answers and summarization.

Dropbox recently acquired Nira, a content governance platform that provides a suite of in-depth functionality to protect cloud files from unauthorized access. With Nira, Dropbox designed a custom solution built directly into Dash, so businesses can easily protect confidential documents in just a few clicks.

Admins can now see everything that's been shared in their company, across every major content platform, in one place. Then, they can identify sensitive content, and manage bulk changes for any number of assets at once - eliminating a previously tedious and manual document-by-document process.

Dash for Business will use self-hosted AI by default, ensuring that customer data remains within Dropbox's trust boundary, without reliance on third-party AI platforms.

https://dash.dropbox.com/

Hyperscience Updates Core Platform

Hyperscience has announced updates to its core platform, the Hyperscience Hypercell, designed to accelerate automation for a wide variety of back-office documents, use cases, and processes.

"The latest release of the Hyperscience Hypercell delivers on the promise of transformation by setting the technology foundation to harness the power of GenAI and LLMs, and enabling our customers to convert back office documents and processes into strategic advantage," said Andrew Joiner, CEO of Hyperscience.

The new version of the Hypercell platform offers new innovations in models, model management, workflow orchestration, and cloud and on-premises infrastructure.

The updates include:

■**New Long-Form Extraction Model -** Long-form documents, like contracts, insurance policies, credit agreements, share purchase agreements, and loan applications, are the lifeblood of key processes within organizations. Traditionally, unlocking critical details from these documents required costly experts who could interpret the nuanced context, spot connections between related information, and extract insights buried within varied formats. The latest version of the Hypercell provides new capabilities to its deep learning model that enable Long Form Extraction and automates this understanding, allowing enterprises to streamline decisions and operations by extracting critical, interdependent data points and multiple occurrences from the complex, unstructured content.

■**Streamline AI / ML Model Lifecycle Management -** As AI and machine learning systems become increasingly integral to daily business operations, effective model lifecycle management is crucial for maintaining seamless execution and compliance with regulatory standards. The complexity of managing model upgrades, governance, and ongoing adaptability often results in resource-intensive processes that slow down innovation and disrupt business workflows. Hyperscience addresses these challenges with new features that help organizations preserve their investment in previously trained models by enabling seamless model portability across different versions of the Hypercell platform.

■**Enhanced Audit Logs -** With the latest release of the Hyperscience Hypercell, organizations have access to enhanced audit logging capabilities that will help users understand everything that happened – who performed the action (the machine or a human), when, and where the action took place. This helps ensure accountability, accuracy, and compliance, thereby enhancing operational transparency and efficiency related to all workflows created within Hyperscience.

https://www.hyperscience.com/

Iron Mountain expands SaaS Solutions

Iron Mountain has announced three new (SaaS) solutions built on its recently launched InSight Digital Experience Platform (DXP), targeting HR, Invoice Processing and Digital Auto Lending.

The Digital HR solution aims to tackle key challenges faced by HR departments by providing secure employee file management in a centralized platform so that both physical and digital documentation is complete, up-to-date, and in compliance with employee records requirements.

It solves common HR challenges by offering a customizable solution with pre-built connectors, workflows, document types, metadata, and AI prompts, all specifically designed for human resources.

Edward Greene, Executive Vice President, Chief Human Resources Officer at Iron Mountain, said: "With Iron Mountain's Digital HR solution, customers can log in to our secure platform for a holistic view of their HR documents instantaneously, regardless of document format.

"We are implementing the solution at Iron Mountain and look forward to enjoying cost savings while enabling our HR team to focus on more strategic initiatives and better serve our employees, customers, and shareholders."

The new digital invoice processing solution built on Insight DXP provides automated data extraction, validation, and matching. It is designed to eliminate manual, slow, error-prone tasks associated with the receipt and processing of supplier invoices.

It offers a customizable solution also with pre-built connectors, automated workflows, document types, metadata, and AI prompts all specifically designed for invoice processing.

The Digital Auto Lending solution for car financiers leverages automation and digitization to transform the auto loan funding process.

Financiers can reduce costs and improve efficiency with intelligent automation, giving them more visibility into the funding process. The solution integrates document ingestion, automated classification, and workflow validation to reduce the time, cost, and degree of manual intervention required to fund an auto loan.

After data extraction, InSight DXP's rules engine analyzes data in the auto finance package for compliance with the bank's loan submission standards. Documents that do not meet these standards are routed to an exception queue for further review, reducing costs, increasing efficiency, and minimizing errors.

Jotform Launches No-Code Automation

Jotform Workflows is a new automation solution based on what was formerly Jotform Approvals, introducing a suite of features that enable users to automate entire processes without writing a single line of code.

With Jotform Workflows, organizations can now receive payments, request signatures, integrate their favourite apps, set up advanced approval systems, and automate a variety of other tasks — all in an easy-to-use, visual workflow builder.

"Jotform Workflows bridges the gap between traditional workflow management and business process management software," said Aytekin Tank, CEO of Jotform.

Key features of Jotform Workflows include

Turn forms into flows: Connect multiple forms, trigger actions with conditional logic and build advanced workflows in minutes.

Streamline task assignments: Assign tasks and manage progress easily, without the hassle of manual follow-ups. Automatically notify clients or team members when tasks are assigned, track progress from anywhere and ensure that every step is completed on time.

Integrate with apps: Connect your workflows with the tools you already use — including Google Drive, Slack, Airtable and more. Automate data sharing to streamline collaboration and keep your team in sync.

Request and receive payments: Add payment requests to any stage of your workflow. Get paid through popular platforms like Square, Stripe, Authorize.net and PayPal — while paying no extra transaction fees to Jotform.

Automate approval processes: Simplify decision-making by automating group and individual approvals. No more back-and-forth emails — team members are automatically notified when their input is needed, and progress is easy to track in Jotform Inbox or Jotform Tables.

Jotform Workflows comes at no additional cost to Jotform users. Enterprise customers also have access to exclusive features, including unlimited workflow runs, white-labelling options, and more.

https://www.jotform.com/products/workflows/

Kodak Alaris integrates GenAI With IDP

Kodak Alaris has announced the integration of generative AI with its Intelligent Document Processing (IDP) software, KODAK Info Input Solution. This new capability makes it easy for organizations to delegate more work to trusted AI services to automate complex document processing tasks with improved speed and accuracy.

With Info Input 7.1, organizations can make better decisions from better data - from any document, in any context, to deliver better business outcomes, faster.

KODAK Info Input Solution is an end-to-end IDP platform built around a unique Open Intelligence™ design, which enables easy integration with AI services from industry giants including AWS, Google, and Microsoft.

Open Intelligence offers faster time-to-value by leveraging pre-built specialized models for the most common document types, such as IDs, claim forms, utility bills, and invoices, and custom models specific to each customer's unique document types.

Global research firm Quocirca noted that Kodak Alaris stands out in the IDP market by combining its capture and information management strengths with its unique Open Intelligence approach to IDP.

According to Jim Forger, VP of Business Solutions at Kodak Alaris, the Open Intelligence approach to AI provides a fast and technology-agnostic path to IDP success, enabling organizations to benefit from the industry's most trusted document AI services that offer a wide range of options, including machine learning, large language models (LLMs), and generative AI.

With Info Input 7.1, Amazon Textract's query capabilities are available directly within the software, empowering users to query data within unstructured documents, ask important questions, and make informed decisions, such as prioritization and workflow routing - before the data even leaves the IDP platform.

In addition, users can leverage Microsoft's OpenAI engine to quickly summarize information within complex unstructured documents and provide context for users to understand the relevance of the data, enabling more informed decision-making.

"The use of LLMs and generative AI significantly improves document processing by automating complex tasks with increased efficiency and accuracy," Forger said.

"The latest release of Info Input Solution enables businesses to extract, validate, summarize, and query data in lengthy unstructured documents such as legal contracts, loan applications, and claim letters, and then prompt the system to deliver contextual information about the data - for example, flagging suspicious responses from a claim submission or delivering a brief document summary and determining what information needs to be redacted. Better data enables better decisions, resulting in better outcomes."

Kodak Alaris has been recognized for its solutions, strategies, and capabilities in the IDP market, earning two Buyers Lab (BLI) Awards from Keypoint Intelligence, a 2024 Fall Pick Award for Info Input Solution as an Outstanding Intelligent Document Processing Solution, and a 2024-2025 Pacesetter Award for Excellence as a Capture & IDP Partner. Analyst house IDC also named Kodak Alaris a 'Major Player' in the IDC MarketScape: Worldwide Intelligent Document Processing (IDP) Software 2023-2024 Vendor Assessment.

Brent Wesler, VP of Strategic Technologies and Digital Automation at PiF Technologies, stated, "Our clients have been asking about how IDP and generative AI will work together, and with Info Input 7.1, I now have an answer for them.

"They can go well beyond extraction and validation, automate more upstream decision-making, and deliver better outcomes to their customers. This is what leading-edge IDP looks like."

Heather Galinis, Accounts Payable Manager for Henley Enterprises, Inc., agrees: "Info Input has revolutionized our key document processes, from AP invoices to COI and tax forms, but it's been difficult to find the right generative AI tool to trust for deeper automation.

"Now that generative AI is part of Info Input - our trusted IDP platform - I'm excited to see how much we can increase our automation and our margins."

For more information, visit https://www.alarisworld.com/ If you are a business in Australia or New Zealand looking to learn more about Kodak Alaris' intelligent document processing solutions, contact the Kodak Alaris Australia & New Zealand Team at Email: service-anz@kodakslaris.com or Dial Toll Free: 1300 252 747.

KnowledgeLake works in Synthetic Labor

Synthetic Labor is a term used by KnowledgeLake to describe a new AI-powered platform that merges Intelligent Document Processing (IDP), Workflow Automation, RPA, and Content Management.

Designed to function as a unified, scalable AI-enabled platform, KnowledgeLake's Synthetic Labor aims to provide an augmentation to human skills and effort.

"With Synthetic Labor, we're not just automating tasks; we're transforming entire business processes," said Russ Malz, VP of Sales for KnowledgeLake.

"This innovation enables organizations to minimize human error, optimize operations, and focus more on strategic initiatives."

The company says the future of work is a hybrid model where Synthetic Labor functions as an intelligent workforce, managing repetitive, high-volume tasks that burden human employees, allowing humans to focus on creative, strategic, and impactful work.

The AI-enabled platform optimizes work distribution, automates low-value tasks, reduces operational costs, ensures consistency, and enhances human collaboration.

By efficiently indexing, categorizing, and routing content, it provides immediate, secure access to vital information across all business applications, empowering organizations to make more informed decisions.

https://www.knowledgelake.com/

Data Discovery for Google Workspace

A tool to discover, classify and secure sensitive data at scale across Google Workspaces has been unveiled by Metomic, a data security and data loss prevention (DLP) solution.

The AI-powered tool automates complex data management Workflows, enabling IT and security teams to maintain control of their data while also ensuring data compliance across cloud-based storage and SaaS environments.

With Metomic's Data Classification solution, organizations can automate complex data workflows and implement "data rules" that ensure files are labelled appropriately within Google.

It also makes it possible to create effective security controls that keep a business' most sensitive data safe from becoming a data security risk (e.g. revoking public access to files marked 'confidential').

Advanced classifying and tagging functionality is a key benefit for highly regulated industries - particularly healthcare and financial organizations - making it easy to adhere to strict regulations such as ISO27001, GDPR, CCPA, SOX, and HIPAA and reducing the risk of costly penalties associated with non-compliance.

Last year, McKinsey & Company surveyed global organizations with more than $US100 million in annual revenue on their data management strategies.

The survey findings revealed 82% of respondents reported spending one or more days per week resolving data quality issues and an astounding 66% were using manual review processes to "assess, monitor, and manage" the quality of their data.

The findings underscore a critical need for innovative, automated solutions that enable companies to identify, classify, label - and secure - their data at scale.

By leveraging its proprietary data loss prevention technology, Metomic's data classification solution allows security leaders to easily add, remove, and modify context-aware data labels across an entire data environment, ensuring data management classification protocols remain accurate and updated.

Not only does Metomic make it possible to automate crucial data management tasks at scale, it enables security leaders to protect the company's most sensitive data against emerging data security risks while maintaining compliance with evolving regulations.

https://www.metomic.io/solution/data-classification

Nasuni integrates M365 Copilot

Nasuni has announced new integration that provides access to the Nasuni File Data Platform via Microsoft Search and Microsoft 365 Copilot, significantly expanding data access for Microsoft's AI services, the company says.

The Microsoft File Share Graph Connector, which can now be leveraged by Nasuni's hybrid cloud platform, is designed to make Nasuni file data easily accessible through Microsoft 365, allowing users to harness the power of their existing office applications while avoiding the pitfalls of AI sprawl and tool overload

Organizations can use the Graph Connector to leverage Nasuni data with Microsoft 365 semantic index to unlock AI-powered search, enhanced compliance, and advanced analytics through a unified interface.

Key benefits of this integration include:

Unlocking Greater Value from Data: Customers can now maximize the value of their Nasuni managed data by making it accessible for personalized experiences via Microsoft 365 Copilot and Microsoft Search, enriching user interactions with relevant content.

AI-Powered Search and Insights: The Graph Connector enables Nasuni managed files to be indexed into Microsoft's semantic index, which forms a key part of the Microsoft 365 AI infrastructure. This semantic index is utilized by Microsoft 365 Copilot and Microsoft Search to provide contextually relevant answers and insights across Microsoft 365 applications.

A Unified Data Interface: Users benefit from single-pane-of-glass access to their Microsoft 365 data (including SharePoint and OneDrive) and Nasuni. This unified view allows for efficient searching and interaction with documents across the entire unstructured file stack, inclusive of Nasuni-managed data.

https://www.nasuni.com/

Onymos DocKnow integrates LLMs

Onymos, developer of solutions transforming Software-as-a-Service (SaaS) for software and application development, has announced the release of an enhanced version of its intelligent document processing component, DocKnow.

The latest version offers a new ability to integrate customer-specific large language models (LLMs), enabling enterprises to extract, process, and validate data from documents with precision and speed.

Onymos DocKnow eliminates the need for time-intensive and error-prone manual data processing by using enhanced optical character recognition (OCR) to extract information from both structured and unstructured data.

This includes printed and handwritten text, numbers, dates, checkboxes, barcodes, QR codes, and more from any document, including personal identification, intake forms, and health and immunization records. DockKnow can also be easily integrated with any third-party back-end information management system – such as Salesforce, AWS, Azure, and Google – or health record system.

In this latest version, DocKnow is strengthened by:

- **A new customer-specific LLM API**: This new API enables enterprises to train their own LLMs using their specific data, resulting in more accurate and domain-specific document processing. For instance, DocKnow reliably and instantly identifies inconsistent data across hundreds of pages.
- **A new, helpful AI assistant:** "Doc," the Onymos AI agent, enables document processing teams – which could include healthcare professionals, legal teams, university registrars, and more – to search through specific documents and hundreds of pages for immediate access to particular information and records.
- **An upgraded, customizable user interface (UI):** The new, simple UI includes bounding boxes, automatic zoom-in/zoom-out, image enhancement, and skew correction, which dramatically improves readability for human reviewers. It allows full customization to match an enterprise's brand, required functionality, and back-end systems. This gives enterprise software engineering and IT teams the ability to modify the component to meet their specific needs as if they had built it from the ground up themselves.

"We understand that many enterprises struggle with time-consuming and error-prone processes like document entry, validation, and retrieval, whether it's for patient care, student registration, or case file review. While these enterprises have started integrating AI tools powered by LLMs like ChatGPT to help with these activities, they often encounter hallucinations and outdated training data issues," said Shiva Nathan, Founder and CEO of Onymos.

"Our enhanced DocKnow addresses these challenges by streamlining document processing and empowering enterprises to train their own LLM models tailored to their specific needs, all while ensuring privacy and security."

As with all Onymos software components, DocKnow is designed with a no-data architecture. This means that all data passing through the solution and used to train the LLM remains securely with the enterprise using the API – no bit or byte of data flows through any Onymos systems or clouds.

Download the white paper here.

https://onymos.com/

Rossum unveils Aurora 1.5 AI Engine

Rossum has announced an update to its intelligent document processing (IDP) solution Aurora 1.5, along with the release of a new Rossum Copilot, which allows users to define data transformations using natural language commands. This update of the Rossum Aurora AI engine brings enhanced features, including a 25% increase in the accuracy of data extracted from documents and faster processing, and includes:

- **Instant Learning for 276 Languages and Handwriting Support***:* Rossum Aurora 1.5 expands the platform's instant learning capabilities to now support 270+ languages, including documents with handwriting.
- **4X Faster Processing for Documents Over 100 Pages***:* Rossum Aurora 1.5 can now process documents with more than 100 pages 4 times faster than before, while maintaining high accuracy in document splitting and data extraction.

The release also introduces the all-new Rossum Copilot, which allows users to define data transformations using natural language commands.

For example, users can:

- Calculate the price after a discount for orders above a certain amount
- Ensure the Document ID is 10 digits by adding leading zeros
- Make sure BIC/SWIFT codes are uppercase
- Remove spaces from IBAN

"Custom Formula Fields were a much-needed feature. Now, we can do things way easier, and instead of maintaining Extensions, we can use Field Manager to align formulas between queues," commented Krzysztof Wis, Automation Lead at IAG GBS (part of International Airlines Group).

"Rossum Aurora 1.5 & Copilot marks a quantum leap in what our AI can achieve for end-to-end document automation," said Tomas Gogar, CEO of Rossum.

"With this release, we further enhance our transactional LLM, empowering global enterprises to scale their automation efforts across new languages, regardless of document volume."

https://rossum.ai/aurora-advanced-ai/

SAP extends GenAI copilot Joule

At its annual SAP TechEd conference, SAP has announced new capabilities that complement and extend Joule, including collaborative AI agents imbued with custom skills to complete complex cross-disciplinary tasks.

Other innovations include the SAP Knowledge Graph, a next-generation solution poised to help developers unlock the full value of SAP data by connecting it with rich business context, and new tools to ensure developers can continue driving Business AI innovation.

On the eve of its first birthday, Joule now features collaborative AI agents to tackle specific tasks and enable them to collaborate on intricate business workflows, adapting their strategies to meet shared objectives.

SAP is infusing Joule with multiple collaborative AI agents that will combine their unique expertise across business functions to collaboratively accomplish complex workflows.

These AI agents enhance productivity by breaking down silos and freeing workers to concentrate on areas where human ingenuity thrives.

Two use cases debuted at TechEd were:

- A dispute management use case employs autonomous AI agents to analyze and resolve dispute resolution scenarios including incorrect and missing invoices, unapplied credits and denied or duplicate payments.
- A financial accounting use case employs autonomous AI agents to streamline key financial processes by automating bill payments, invoice processing, and ledger updates while quickly addressing inconsistencies or errors.

A new SAP Knowledge Graph solution, accessible through SAP Datasphere and Joule in Q1 2025, will give users a deeper layer of business understanding by seamlessly mapping relationships and context across SAP's vast data landscape, empowering organizations to make better decisions with their data.

By offering ready-to-use relationships between business entities like purchase orders, invoices, and customers, the solution can significantly reduce the complexity of manual data modelling.

SAP Knowledge Graph grounds AI in SAP-specific business semantics, which reduces the risk of inaccurate or irrelevant results and makes it easier for organizations to build intelligent applications and leverage generative AI more effectively.

SAP also launched new generative AI developer capabilities such as code explanation and documentation search in SAP Build, the company's platform for extending its solutions, which will reduce development time for Java and JavaScript developers.

SAP Build is also adding an Extensibility Wizard feature that will let developers access SAP Build directly from SAP S/4HANA Cloud Public Edition, simplifying the extension process. Meanwhile, ABAP developers and fusion teams will get seamless access to ABAP Cloud development tools from SAP Build.

Reltio enhances Purview Governance

Reltio has introduced Reltio Integration for Microsoft Purview, which combines Reltio's data unification and management capabilities with Microsoft Purview Data Governance.

Reltio Integration for Microsoft Purview makes it easy to discover and consume trusted core data and 360 customer profiles in enterprise-wide business initiatives. It also enables greater visibility into and governance across the enterprise's full range of data assets.

"The explosion in data and the number of applications enterprises maintain has created great demand for reliable data and simplified data governance. By better managing data assets and ensuring effective data discovery, lineage, quality, and compliance, enterprises can accelerate the success of their business initiatives and make better data-driven decisions," said Venki Subramanian, Senior Vice President of Product Management at Reltio.

"By integrating with solutions like Microsoft Purview, we are making it easy and cost-effective for enterprises to transform their fragmented data into unified data while also moving to more federated data governance."

"The integration with Reltio enables our customers to create unified, trusted data in Reltio, publish them, and associate them with Master Data Products in Purview for further curation and governance," said Karthik Ravindran, General Manager, Enterprise Data at Microsoft. "This provides customers with an integrated MDM experience in Purview with Reltio, allowing them to effectively govern their data estates."

Features include:

- **Data discovery**: More easily and fully manage data with simple searchability, easy organization, and at-a-glance visibility of data quality and health.
- **Trusted core data and customer 360**: Unify, standardize, and enrich core data with Reltio Multidomain MDM and deliver trusted 360-degree customer views across the enterprise with Reltio Customer 360 Data Product.
- **AI innovation**: LLM-powered entity resolution, automatic anomaly detection, and genAI-powered data exploration, segmentation, and chat-based search.

Reltio is available on the Microsoft Azure Marketplace.