



Quarterly Newsletter

February 2026

Inside This Issue

- PG. 2**
President's greetings
- PG. 3**
EU AI Act “moves from theory to operating model” (and may still get tweaked)
Canada’s AI regulation path: AIDA is stalled, but governance expectations keep rising
- PG. 4**
ISO/IEC 42001 takes a big step: AI governance becomes certifiable (for real)
- PG. 5**
AI security & agent governance: NIST guidance (models, prompts, and data become attack surfaces)
Content authenticity & provenance: C2PA/Content Credentials trend to “data governance for media”
- PG. 6**
Privacy-preserving data collaboration is going mainstream
Metadata/catalogs become the “control plane” for AI (RAG + agents need trusted context)
- PG. 7**
Closing remarks and how to stay involved



Welcome to our first DAMA Edmonton newsletter of 2026! Data management is moving fast—AI governance, new legislation and directives, GenAI, composable architectures, and AI agents are changing what “good” looks like in every industry. Thanks to our members and partners for keeping this community strong. This year we’ll keep bringing in thought leaders, creating high-value networking at chapter meetings, running workshops shaped by member feedback, and adding more hands-on learning opportunities by hosting events on timely problems that matter to Alberta’s data community.

We’re also entering an exciting leadership transition. After six years as President, I’m proud of the growth and success we’ve achieved together, and grateful for your support throughout that journey. I’m pleased to welcome **Peter Guest** as our **incoming President**. We’re also excited to add **Bhavya (Bee) Jha** as Director, Marketing & Communications. With appreciation for their contributions, we acknowledge the resignations of **Surinder Padem** (VP, Marketing & Logistics), **Eric Dressler** (VP, Data Science Initiatives), and **Mona Elyas** (Associate Director).

Looking ahead, please mark your calendar for **the 12th Annual DAMA Edmonton Data Management Conference** (October 2026) at **the NAIT Productivity & Innovation Centre (10210 Princess Elizabeth Avenue NW)**—a two-day main event with workshops before and after, plus certification exams following the conference. And don’t miss our **monthly chapter meetings**, held **the last Thursday of every month, 5:00–7:00 PM at Edmonton Unlimited (10107 Jasper Avenue NW)**. We’ll share upcoming speakers and roundtable topics in future updates and at www.dama-edmonton.org.

Warm regards,

Stephen Kaiser
President, DAMA Edmonton Chapter

EU AI ACT “MOVES FROM THEORY TO OPERATING MODEL” (AND MAY STILL GET TWEAKED)

Europe’s AI “rulebook” has officially left the policy shelves and entered the day-to-day operating model. The EU AI Act is rolling out in waves—starting with baseline provisions (including AI literacy and certain bans) and moving toward obligations on general-purpose AI (think foundation models) before a full roll-out by August 2027. ([AI Act Service Desk](#))

For data leaders, the message is simple: regulators want “audit-ready” evidence. Not just what your model does, but what data shaped it, how it was tested, and what controls keep it safe in production. The EU’s voluntary GPAI Code of Practice is a starter kit—practical expectations and templates for transparency, copyright, and safety/security that many firms will treat as the benchmark. ([Digital Strategy](#))

Why should Canada care? Because global organizations don’t build one governance system per region—they standardize, and the EU is often the strictest customer at the table. And the EU may still tweak details: the Commission’s “Digital Omnibus on AI” proposes targeted simplifications to parts of the Act amid industry pressure. Use this moment as a forcing function: tighten documentation, lineage, and accountability now—then reuse those capabilities everywhere. Next: certifiable AI management systems. ([Digital Strategy](#))

CANADA’S AI REGULATION PATH: AIDA IS STALLED, BUT GOVERNANCE EXPECTATIONS KEEP RISING

While Europe is putting the EU AI Act on a rollout calendar, Canada’s national AI rulebook is still being drafted in the margins. Bill C-27—home to the proposed Artificial Intelligence and Data Act (AIDA)—died on the Order Paper when Parliament was prorogued on January 6, 2025, and it hasn’t returned as a single, clear “AI law” yet. ([parl.ca](#))

But “no AIDA (for now)” doesn’t mean “no expectations.” Privacy reform, regulators, and major customers are steadily raising the bar on proof: what data you used, how you tested, who is accountable, and how you detect and respond when an AI system misbehaves. Think of it as governance by gravity—standards and oversight pull everyone toward safer practice. ([Osler, Hoskin & Harcourt LLP](#))

The smart business move is to build the capabilities you’ll need anyway: a model inventory (what’s deployed where), data lineage (where it came from), risk assessments, logging, and retention. ISED’s AIDA materials outline the direction of travel—and firms that get “audit-ready” now will sell faster, partner easier, and be ready when Canada’s next legislative wave arrives. Next up: making AI governance certifiable with ISO/IEC 42001. ([ISED Canada](#))

ISO/IEC 42001 TAKES A BIG STEP: AI GOVERNANCE BECOMES CERTIFIABLE (FOR REAL)



Image source: <https://www.forbes.com/sites/bernardmarr/2025/02/03/generative-ai-vs-agentic-ai-the-key-differences-everyone-needs-to-know/>

AI governance just got a lot less “hand-wavy.” In January 2026, UKAS granted its first accreditation to certify organizations against ISO/IEC 42001—meaning “responsible AI” can now be audited and certified, not merely promised in slide decks. For businesses, that’s the start of an AI trust badge customers and regulators can recognize. ([UKAS](#))

ISO/IEC 42001 is essentially ISO 27001’s “management system” idea—applied to AI: clear roles, lifecycle controls, measurable objectives, incident response, and supplier oversight. Accreditation bodies like ANAB are already standing up programs for AI Management Systems, turning policies into evidence you can show during procurement, partner due diligence, or an internal audit. ([ANAB](#))

Why should Canada care? Because global clients increasingly ask, “Prove your AI is governed.” Canada’s Standards Council has an Artificial Intelligence Management Systems accreditation scheme emerging as well—so the play is to build reusable governance now: model inventory, lineage, monitoring, and documented decision rights. Next, we’ll look at how trust also shows up in data sharing—via privacy-preserving collaboration and clean rooms. ([scc-ccn.ca](#))

AI SECURITY & AGENT GOVERNANCE: NIST GUIDANCE (MODELS, PROMPTS, AND DATA BECOME ATTACK SURFACES)

After provenance comes protection: AI is becoming the newest “employee” in many firms—fast, helpful, and occasionally gullible. The more we let AI read, write, and act, the more it becomes part of our cyber perimeter. In late 2025, NIST released draft guidance and a Cybersecurity Framework Profile for AI to bring these risks into a familiar enterprise playbook. ([NIST](#))

For data teams, the takeaway is refreshingly practical. Secure not only models, but also the data that shapes and feeds them: training sets, retrieval libraries for RAG, prompts, logs, and the connectors that let tools touch production systems. CISA frames this as protecting the “AI data supply chain,” with controls you can actually evidence in audits and vendor reviews. ([CISA](#))

The newest wrinkle is agentic AI—systems that can plan steps and call tools. NIST’s CAISI has opened an RFI on securing AI agent systems, focusing on risks that appear when model output can trigger real actions. Businesses that bake in least-privilege permissions, approvals for high-impact actions, and end-to-end logging will ship faster—and earn trust across borders, including Canada-EU partnerships. ([NIST](#))

CONTENT AUTHENTICITY & PROVENANCE: C2PA/CONTENT CREDENTIALS TREND TO “DATA GOVERNANCE FOR MEDIA”

In 2025, “AI slop” flooded our feeds; in 2026, brands are starting to ask a sharper question: can we prove what’s real? The Content Authenticity Initiative argues we’ve hit a turning point—Content Credentials are moving from a nice idea to something created at capture, carried through workflows, and increasingly understood by end users. ([contentauthenticity.org](#))

The engine behind this is C2PA, a kind of tamper-evident “nutrition label” for media: it attaches provenance metadata and a cryptographic signature that can record key edits and help confirm an asset’s origin. For business, that’s practical risk reduction—fewer fake endorsements, clearer rights management, stronger evidence for investigations/claims, and faster decisions when “Is this authentic?” suddenly becomes a board-level question. ([C2PA](#))

Platforms are even floating a strategy shift: fingerprint authentic media rather than endlessly chasing fakes—pushing provenance standards into the mainstream. Canadian organizations can get ahead by requiring Content Credentials in agency contracts, capturing provenance at creation, and training teams to verify before sharing. Next, we’ll zoom out: trust isn’t only about media—it’s also about securing the AI systems that consume it. ([TechRadar](#))

PRIVACY-PRESERVING COLLABORATION GOES MAINSTREAM: CLEAN ROOMS + PET CONTROLS MATURE



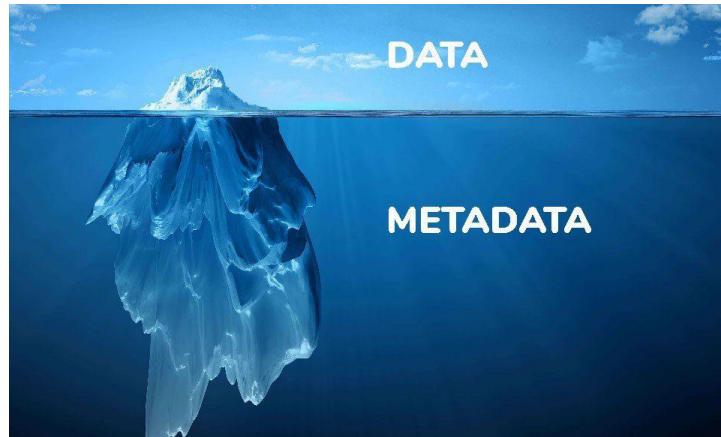
Source: <https://cdp.com/articles/data-clean-room-data-privacy/>

Data clean rooms are the “glass-walled boardrooms” of the data world: two (or more) organizations can learn something together—campaign lift, churn drivers, fraud patterns—without handing over raw customer records. Built on privacy-enhancing techniques (PET) and strict governance, they’re becoming the polite way to collaborate when lawyers (and customers) say, “no raw data sharing.”

What’s changed recently is how mature the guardrails have become. AWS added cross-region collaboration (so partners don’t need awkward data copies) and “data access budgets” that cap how often shared tables can be queried or used for ML—turning privacy from a promise into a measurable control. ([Amazon Web Services, Inc.](#))

For Canadian organizations, this is a competitive lever with global reach: you can work with EU/US partners, unlock joint analytics, and still keep audit logs, approved queries, and output controls that reduce re-identification risk. With vendors like Snowflake continuing to ship clean-room updates, this is moving from niche to normal—and it tees up our next topic: metadata and lineage as the navigation system for governed sharing. ([docs.snowflake.com](#))

METADATA/CATALOGS BECOME THE “CONTROL PLANE” FOR AI (RAG + AGENTS NEED TRUSTED CONTEXT)



Source: <https://zeenea.com/the-role-of-metadata-in-a-data-driven-strategy/>

After you’ve learned to collaborate safely (clean rooms PETs), the next problem is deceptively human: “Which data should we trust?” That’s why metadata catalogs are suddenly the AI era’s control tower. Snowflake’s decision to bring Select Star into Horizon Catalog is a clear bet that better context—definitions, ownership lineage—makes AI answers safer and more useful ([InfoWorld](#))

Catalogs are also graduating from “inventory lists” to active governance systems. Microsoft Fabric made its Govern experience generally available and exposed governance capabilities through public APIs—signaling that policy, security, and reuse are becoming built-in platform features, not afterthoughts. ([Microsoft Fabric Blog](#))

For Canadian organizations, this is a growth lever: faster AI delivery, fewer stalled projects, and cleaner audits because everyone can see what data means and where it came from. Azure Databricks’ recent update to sharing Unity Catalog assets reinforce the same trend—treat data and AI assets like permissioned products you can safely reuse. Next, we’ll take this “trust layer” beyond data: proving authenticity and provenance in a world of AI-generated content ([learn.microsoft.com](#))

CLOSING REMARKS AND HOW TO STAY INVOLVED



Image source: <https://www.dama.org/cpages/chapters>

As we kick off 2026, one theme keeps repeating: the organizations that win with AI and data are the ones that treat data management as a community sport. DAMA Edmonton is your local “home base” for practical governance, modern architecture, and real connections—whether you’re a seasoned professional or a student trying to break in.

The easiest way to contribute is to show up and bring one colleague. Join our chapter meetings (networking included), share upcoming events with your team, and help amplify speakers and topics on LinkedIn. If you want a bigger role, volunteer—event logistics, registration, sponsorship support, communications, and content all make a measurable difference (and look great on a résumé).

Our monthly chapter meetings run the last Thursday of every month, 5:00–7:00 PM at **Edmonton Unlimited, 10107 Jasper Avenue NW**. Last month’s meeting was on January 29, featuring **Russ Greiner** (U of A Computing Science; founding Scientific Director of Amii; CIFAR AI Chair) on “**Towards Patient-Specific Treatment: Medical Applications of Machine Learning.**” Register for the February 26 meeting with **Sam Jenkins, Co-founder & Managing Partner of Punchcard Systems**; CEO of Standard Field Systems, Co-founder, Wellnext and other companies - a Serial entrepreneur, Chair Board of Directors, Winspear, Past Chair Fringe Festival topic **From Data to Decisions: What High-Performing Digital Organizations Do Differently** via our website <https://dama-edmonton.org> And start planning now for the 12th Annual DAMA Edmonton Data Management Conference (Nov 4-7, 2026) at the **NAIT Productivity Innovation Centre**—with workshops before/after and certification exams.