

Issue No. 27, Spring 2026

Think TWENTY20

The Magazine for Financial Professionals



The Rise of AI in Accounting – Interview of Liv Watson; Ethics, Technology and the Future of the Profession; Standardization Meets Innovation; How AI and XBRL Are Reshaping Financial Analysis; AND The Fusion of Risk Management for Artificial Intelligence and Human Agents

Number 27, Spring 2026

Editor in Chief: Gerald Trites, Managing Editor: Gundi Jeffrey, Contributing Editor: Eric E Cohen

Email: info@thinktwenty30.com Telephone: (416) 602-3931

We are a community of people who want to learn about, and discuss, important issues. To join us, please go to our website at <https://www.thinktwenty20.com>. Regular subscription – free. Premium subscription, One year - \$30, Two years - \$50.

ISSN 2563-0113

Cover picture; thanks to chatgpt.

FOUNDING PARTNER



Supporter



University of Waterloo Centre for Information Systems Security and Assurance

Table of Contents

Editorial.....Page 3

By Gerald Trites, FCPA, FCA, Editor-in-Chief

The Rise of AI in Accounting: In Their Own Words.....Page 4

By Gundi Jeffrey

As the global AI-in-accounting market continues to grow at a projected CAGR (Compound Annual Growth Rate) of over 41% through 2029, the technology is becoming a cornerstone of modern financial operations.

Ethics, Technology and the Future of the Profession.....Page 13

By David Wray, ACA, CPA

A growing tension is emerging between technology development, particularly for artificial intelligence (AI), intellectual property and data control rights and sovereignty.

Standardization Meets Innovation: AI Platform for Regulatory Data...Page 21

By Björn Fastabend

Rushing into AI adoption without a coherent strategy easily leads to a "zoo" of disconnected applications, each with its own requirements, maintenance overhead and integration challenges. How can this be avoided?

Beyond Compliance: How AI and XBRL Are Reshaping Financial

Analysis.....Page 26

By Janis Steinmann

The combination of structured data and AI enables analysis that was previously impractical at scale. Financial reporting will increasingly look less like documents and more like data services: verifiable, continuously updated, with enhanced assurance built in.

The Fusion of Risk Management for Artificial Intelligence and Human

Agents.....Page 33

By Eric E. Cohen, CPA

This article describes a collaborative effort to create something called the *Agent Metadata Specification*. The work is an open-source project, with a goal of supporting consistent and more formalized agent risk management.



***Gerald Trites, FCPA, FCA
Editor-in-Chief***

Our cover this issue shows a human reacting despairingly to being checkmated by a robot. Of course, computers have been beating chess players, even world champions, ever since Deep Blue beat Gary Kasparov in 1996. And AI has advanced tremendously over the past thirty years.

In 2022, Open AI introduced Chatgpt, which set a new standard for AI by using large language models – specifically generative pre-trained transformers (GPTs) – to generate text, speech and images in response to user prompts. It greatly accelerated the growth of AI. The past few years have been marked by major change, including the advent of Agentic AI, which provides more independent decision making and Identical AI, which is like agentic AI but is focused on enhancing and developing an individual’s digital identity, and promises to be a major innovation in itself. It will likely lead to people developing their own digital identity, which will become a major part of how they present to the world and eventually lead to changes in the meaning of personal identity, perhaps even to changes in what it means to be human. Some experts maintain that AI is nearing the level of artificial general intelligence (AGI) – AI that matches or surpasses human capabilities across virtually all cognitive tasks.

Personal robots already exist, but are very expensive and still have a level of intelligence that could be/will be improved. But they, too, are being enhanced with the new levels of AI now developing and will, no doubt, reach AGI in the not-to-distant future.

We have only just begun to witness the changes to society to be wrought by AI. It will be a new world! This magazine will contain much more content to explain and explore the further development of AI as used in the accounting and financial professions.

The Rise of AI in Accounting: In Their Own Words

By Gundi Jeffrey



Gundi Jeffrey is an award-winning business journalist specializing in writing about the accounting profession for various publications in Canada and England. In 1985, she co-founded The Bottom Line, Canada's only independent publication for the accounting and financial professions at the time, serving as its executive editor.

Artificial intelligence has rapidly shifted from a futuristic concept – or at least one limited to the Big 4 – to a practical, transformative force within the accounting profession. Today, AI has the potential to be reshaping how financial data is processed, analyzed and leveraged for decision-making. Its adoption is accelerating across firms of all sizes – a from global accounting giants to small and mid-sized practices – mostly driven by market forces, competitive demand and the need for greater efficiency. As the global AI-in-accounting market continues to grow at a projected CAGR (Compound Annual Growth Rate) of over 41% through 2029, the technology is becoming a cornerstone of modern financial operations.

Large global firms such as Deloitte, PwC, EY, and KPMG are leading the charge. These organizations are investing heavily in generative AI and intelligent automation to streamline audits, enhance financial reporting, and improve advisory services.

Mid-sized and boutique CPA firms are also increasingly integrating AI tools to stay competitive. Many accountants report using AI for tasks ranging from mail automation to drafting technical memos, data analysis and even marketing.

And, of course, businesses and finance departments across industries are also embracing AI to improve internal accounting workflows.

Despite the massive positive hype surrounding AI, it has its detractors. They worry about systems making decisions they don't understand or that tech companies are gaining too much influence. They are afraid of a future where humans rely on machines for everything. And – along with promoters and detractors, there are those who want to take a balanced and cautious approach – identifying potential benefits, characterizing risks, and determining ways to mitigate those risks.

To help clarify where all this is heading, we spoke with Liv A. Watson, a Senior Fellow at the Data Foundation, contributing to both Special Projects and the Climate Data Collective. She co-founded the Global Digital Single Market Data Alliance (DSMDA) project, as well as the open data standards XBRL and the XBRL International Consortium. Her extensive work has been crucial in the advancement of global digital business reporting standards. When agreeing to this interview, she said that AI in the accounting world is “a topic of passion” for her.



ThinkTWENTY20: *Your quote about AI in accounting being “a topic of passion” immediately caught our attention. Can you tell us a bit about how this became a personal crusade?*

Watson: This goes back to my university days when professors taught us how to copy and paste data into a spreadsheet, and I thought “is this why I’m going to college?” And, as the internet was still in its early days, I started to look at emerging technology as something that would get me out of copying and pasting data into spreadsheets. So, I was thinking to myself that this is the process we can go to apply our expertise. I kept thinking that a machine could learn to do these processes. I then got quite involved with the universities and the thinkers on this topic around the world, and it’s been a passion ever since.

AI adoption is accelerating across firms of all sizes – a from global accounting giants to small and mid-sized practices – mostly driven by market forces, competitive demand and the need for greater efficiency.

ThinkTWENTY20: *Where do you see AI making the most meaningful impact on the day-to-day work of accountants and finance professionals?*

Watson: For me, it always comes back to that instinctive question: Isn’t there a better way to do this? And right now, AI is stepping in exactly where accountants spend most of their time— by handling repetitive tasks, processing huge volumes of data, and making routine judgments. We’re already seeing AI take over things like data entry, invoice processing, reconciliation, and exception handling. These are the kinds of activities that can eat up hours, and AI can now do them quickly, consistently, and at scale. That shift gives accountants more time to actually dig into the data, understand what it means, and deliver insights that matter. On top of that, AI agents can run 24/7, monitoring workflows, suggesting journal entries and performing calculations without ever slowing down. What’s really interesting is how these AI agents are becoming part of the process itself.

My observation is that the Big Four, for example, are building their own specialized agents to support audit and assurance work. We're moving from traditional sampling to the ability to test entire populations of transactions. Something that fundamentally changes what's possible in an audit. So, I think we're entering a phase where AI isn't just improving individual tasks; it's becoming woven into the accountant's daily workflow. It's helping firms rethink how audits are done, how assurance is delivered, and how professionals spend their time. And ultimately, it's elevating the role of the accountant by shifting the focus from processing data to interpreting it and adding real strategic value.



ThinkTWENTY20: *I've also read that AI is quite useful in fraud detection. How do you see AI tools changing how firms identify and deal with fraud?*

Watson: AI can do a lot, right? It can look at large data sets to look at patterns that we humans would have a difficulty to catch. But I also still see limits to that. Take identity fraud. There is the data, but you also need to interview the executives

involved, understand the pertinent internal controls, and the things that go into audit and assurance. Modern fraud is very sophisticated, right? What about fake invoices? What about fake agents, all of those things. I do think we are still in the early stages of this work. We still cannot put our full trust in AI agents for effective fraud detection.

I think we're entering a phase where AI isn't just improving individual tasks; it's becoming woven into the accountant's daily workflow.

Also, think about history or continuous auditing. There are a lot of false positives that get flagged. For example, if we get too many false positives, how is the human going to be able to interpret that? But I do see that AI is going to have a huge impact on fraud analysis, and it will be able to bring in all kinds of data set analyses, and truly change the profession in terms of what we do and how we detect fraud. And, instead of being the ones who come in after everybody is wounded, we are now going to be much more proactive in looking for and detecting fraud. So, I think there will be a lot of opportunities for using AI for fraud detection. But we are still a long way from having 100% trust in the outcome of these agents.

ThinkTWENTY20: *Financial reporting and auditing are still the core of the profession. How is AI changing the nature of both of these disciplines?*

Watson: What we're seeing now is that a lot of the rules and regulations that guide financial reporting are essentially becoming executable code. Software vendors are building those requirements right into the tools companies use every day. But that doesn't replace the need for accuracy, assurance or transparency. If anything, it shifts where the effort goes. Financial reporting used to be all about producing the numbers such as pulling data, reconciling accounts, checking for inconsistencies, drafting the narrative. AI is flipping that workflow.

Instead of spending most of our time creating the report, we're spending more time analyzing what the data is telling us.

I think back to the early days of financial reporting, when companies would send their final numbers to a printer who would design and assemble the report. That "last mile" was always the most stressful part of the cycle – every quarter, or every six months in Canada. Then the cloud came along, and companies like Workiva changed everything by making real-time collaboration and a single source of truth linked to multiple destinations possible. Now they're layering AI on top of that. The result is that closing the books or preparing disclosures becomes far less frantic. We're moving toward near real-time reporting instead of the monthly or quarterly scramble we're all used to.

Auditing is going through a similar shift. As I mentioned earlier, AI makes continuous monitoring possible and allows us to look at 100 percent of transactions instead of relying on sampling. That's a huge change. But even with all that automation, professional judgment doesn't go away. AI changes how we perform the work, not the need for expertise. It gives auditors better tools, better visibility, and more confidence, but I still believe that the human element is still essential.

The challenge, and the opportunity, is figuring out how we develop future accountants so they're not just operating AI tools, but are actually understanding the work behind them.

ThinkTWENTY20: *A recurring theme is that AI is taking over time-intensive, repetitive tasks – transaction coding, reconciliation, data extraction. What are the genuine benefits of letting AI take on the more mundane tasks necessary to get accounting work done?*

Watson: The obvious benefit is that it frees accountants to actually do the work we're trained for. Most of us didn't go to school dreaming about coding invoices or chasing down mismatched entries. When AI handles those routine tasks, it gives professionals more time to focus on analysis, judgment and the bigger-picture thinking that really adds value.

But I also think we need to talk about the flip side. Those repetitive tasks have traditionally been where new accountants cut their teeth. That early-career grind – learning the basics, understanding how transactions flow, getting a feel for the numbers – that's how many of us built our foundation. If AI takes over all of that, we risk losing an important part of the training pipeline. My worry is that the profession could drift toward being "AI-dependent" instead of "AI-intelligent." If we rely too heavily on the technology without making sure people still understand the underlying principles, we could end up with professionals who can operate the tools but don't fully grasp the accounting behind them.

So, yes, there are real benefits with more efficiency, more time for meaningful work, fewer late nights doing manual tasks. But we also have to be intentional about how we train and develop young professionals in an AI-driven environment. The technology should elevate the profession, not hollow out the skills that make accountants trusted experts.

ThinkTWENTY20: *There's also the possibility of job displacement, a concern that is real and often discussed. But, perhaps less examined, is a subtler worry: that over-reliance on AI might erode the analytical and creative thinking skills that define what a good accountant actually is. Do you think that's a valid fear?*

Watson: Yes, I'm worried about that loss and I do think it's a valid fear. I'm genuinely worried about losing those core skills that make accountants good at what we do. My advice to students today is pretty simple: keep learning. Lean into understanding both the technology and the profession and, "most importantly," how the two fit together. But don't hand over your judgment to AI. Critical thinking, skepticism and ethical reasoning are still the heart of this profession, and they're going to matter even more as AI becomes more embedded in our work.



At the same time, job displacement is a real concern. When I was growing up, if someone asked what you wanted to be, saying "an accountant" felt like a solid, respected answer. Today, parents and grandparents are asking a very different question: *What career will still exist in a world where so many roles, "not just in accounting," are being automated or enhanced by AI?* It's a fair question. I absolutely believe accounting still has a place. The profession isn't going away. But we do have to be honest: if AI takes over a lot of the routine, mundane tasks, we may not need as many people doing those tasks. The challenge, and the opportunity, is figuring out how we develop future accountants so they're not just operating AI tools, but actually understanding the work behind them and bringing the judgment, insight and integrity that no machine can replace today.

ThinkTWENTY20: *There's also concern about how AI uses information to produce expected results and whether AI is collecting more information than we might realize. Uncertainty breeds distrust. How do you respond to those concerns?*

Watson: I think that concern is absolutely valid. A lot of what we call “AI agents” today are really more like shadows of agents. What I mean is that we see the output, but we don’t always understand what’s happening behind the scenes. And that’s the biggest challenge right now. When we use an AI agent, how is it functioning? What data is it using? What decisions is it making along the way? That transparency gap is the missing link.

I’m working with a group called TAVRO AI, which is trying to tackle exactly this problem by developing an open metadata standard. The idea is simple but powerful and grounded in the idea that every AI agent should come with basic documentation. Who owns it, what it’s designed to do, what data does it have access to, how it operates and what its capabilities and limitations are. Without that metadata, the agent is basically a black box. With it, we can start to understand how the outputs are produced and what reasoning sits behind them.

This kind of standardization is going to be essential if we want people to trust AI systems. It gives organizations a way to govern their AI, assess risks and make informed decisions about how these tools are used. And we need a global conversation around this, just like we had with XBRL years ago. TAVRO AI built the first version of what this standard could look like, and the next step is to bring regulators and standards bodies into the discussion so it can eventually be embedded into legislation.

Because here’s the reality: companies aren’t just using one or two AI agents anymore. Some of the big players have thousands. And many of them don’t fully understand the risks or even the basic functionality of the agents they’re relying on. That’s not sustainable. If AI is going to be trusted and widely adopted, we need clear rules of the road and metadata is a big part of that.

Criminals adapt quickly, often faster than we do, and they’re already experimenting with ways to use AI to create synthetic transactions, fake evidence and highly convincing financial documents.

ThinkTWENTY20: *There's significant unease about what data AI systems are trained on, what they retain and who has access to it – especially when firms are feeding sensitive client financial data into third-party tools. How should accounting professionals think about data sovereignty and privacy when adopting AI agents?*

Watson: This is one of the biggest questions the profession has to grapple with right now. When we talk about AI, we often jump straight to the cool capabilities, but the real issue is: who controls the data, where does it live and how is it governed? Accounting firms can’t treat data as something that just sits on servers or laptops. It needs to be governed with intention. Data governance is about more than just ensuring data is machine-readable or accessible; it is about establishing precise control over who can access it and defining their authorized actions.

The necessity of this was highlighted during one of our projects when an AI agent created by a developer triggered an immediate security alert, asking individuals to reset their passwords. This served as a critical reminder of how rapidly an agent can interact with systems in unforeseen ways, underscoring why robust governance is essential.

Firms need to minimize exposure. Sensitive financial data shouldn't be floating around in places where you can't track who sees it. You need to know where the data is stored, how it's protected and what happens to it once an AI agent touches it. And this ties back to metadata standards. If we don't know what an AI agent is capable of "what data it uses, what it retains, how it processes information," then we can't trust it.

That's why I'm working with TAVRO AI on an open metadata standard. The goal is to create a baseline: every AI agent should come with clear documentation about what it does, who owns it and how it behaves. Without that, it's a black box. With it, we can start to build transparency and trust. And, honestly, I don't think AI adoption is primarily a technology decision anymore. It's just as much about the governance decision. We're already seeing companies deploy hundreds or even thousands of AI agents, and many don't fully understand the risks. That's not sustainable. We need global dialogue, regulators at the table and clear expectations for documentation and transparency. Data sovereignty and privacy aren't side issues – they're becoming the foundation of how AI will be used responsibly in accounting.

ThinkTWENTY20: *Beyond data privacy, there are harder ethical questions: bias baked into training data, accountability when AI-generated outputs are wrong, and the use of AI to generate convincing but false financial documents or communications. Are these concerns the profession is taking seriously enough, or are they being drowned out by the enthusiasm?*

Watson: I don't think these concerns can be overstated. In fact, I'd say they're some of the most important issues the profession has to wrestle with right now: bias in training data, misuse of AI and the possibility of generating fake invoices or financial statements. These aren't theoretical risks. They're real, and they demand new controls, new audit methodologies and a much stronger data governance framework than we've ever had before. AI is powerful and, like any powerful technology, it can be used for good or for harm. That's why we can't just assume the outputs are correct or unbiased. We still need human judgment in the loop. That is people who can look at an AI-generated result and ask, "Does this make sense? How did the system get here?" Without that, we're flying blind.

And, yes, the enthusiasm around AI is huge, but that can't drown out the ethical questions. If anything, it should make us more determined to address them. We need to understand how these systems work, what data they're trained on and what risks they introduce. That's where governance comes in again. Strong data governance, clear documentation and transparency around how AI systems operate will be essential. So no, these fears aren't exaggerated.

ThinkTWENTY20: *There's a darker possibility that doesn't get discussed enough: that AI tools could be weaponized within the financial system itself – to manufacture misleading reports, game regulatory filings, or enable sophisticated financial fraud at scale. Is that a realistic concern, and what role do standards bodies and regulators have in getting ahead of it?"*

Watson: It's absolutely a realistic concern. AI can be used for criminal activity, and we'd be naïve to think otherwise. Criminals adapt quickly, often faster than we do, and they're already experimenting with ways to use AI to create synthetic transactions, fake evidence and highly convincing financial documents. Traditional controls simply aren't built for this new reality. That's why I keep coming back to transparency and governance. Regulators and standards bodies need to require that the AI agents we use are fully documented and understandable. If we don't know how an agent works, what data it touches or how it generates outputs, then we can't detect when something has been manipulated – or when an agent is being used maliciously.



We also need new policies and frameworks to understand these emerging forms of financial crime. We haven't yet seen the full extent of how AI could be weaponized, but we can already imagine the possibilities. And the worst thing we could do is wait until after something happens to respond. Regulators and standards setters need to be out in front, not playing catch-up. The challenge is that standards often take years to develop and, by the time they're released, the technology has already moved on. That's why we need a more agile approach – one that brings regulators, technologists and the profession together early and often. If we want to stay ahead of AI-enabled fraud, we can't rely on yesterday's rulebook. We need governance frameworks that evolve as quickly as technology does.

ThinkTWENTY20: *How do you see AI reshaping the accounting profession over the next three to five years – and are there particular considerations for Canadian/global firms and the Canadian/US/EU regulatory environment that differ from the global picture?*

Watson: Like I said, AI will become embedded, layered across all workflows. So, I think that is transformative, and we are moving to real-time operations. We used to live as historians of finances and now, finally, we're moving toward continuous assurance and auditing. The academic world has talked about this for a long time, but I think now we're going to see mass adoption in the marketplace of these ideas and technology advancements.

I think we are moving towards being more strategic accountants, not just somebody who can understand the rules and regulations and know a debit from a credit. We have to become more strategic to be able to use the data we have and how to communicate it. As for the regulatory environment, we need to move at the fastest pace possible to get the necessary regulations out there. But, if you look at Canada, or Norway, it's all very much the same as always. They've got sovereignty and transparency on AI in their environment. The US regulators are demanding more evidence and regular documentation of the AI. Meanwhile, the EU leads with the EU AI Act. They want stricter rules, not just documentation and transparency. They want to develop rigorous oversight and enforcement requirements.

ThinkTWENTY20: *Your work on XBRL was fundamentally about making financial data interoperable and machinereadable – a kind of prerequisite infrastructure for what AI is now being asked to do. Looking back, does it feel like the profession built the runway without quite knowing what would eventually take off from it?"*

Watson: I feel like we built the infrastructure while it was being built. It's like flying the airplane as you were designing it. What we thought during the dot.com era was that everybody would adopt XBRL much sooner, because it made data machine-readable and discoverable, and it obviously put structure and governance around data. But it took a long time to get people on this journey. It goes back to the data we enter into AI, or what AI analyzes of that data. The data will be of much better quality if it is structured and has controlled vocabularies, controlled taxonomies.

So, I think XBRL really was the runway without us knowing where it really would take us. Now we have built the roads, and we need to build the cars, the gas stations, and those kinds of things to create the value add. But I don't think that some people say that AI will make XBRL obsolete. Where I do think AI has a huge value right now is that where we previously had to manually, or auto-manually, tag information, we can now automate that tagging. I think creating an XBRL document was a primer for AI advanced solutions.



Ethics, Technology and the Future of the Profession

By David Wray, ACA, CPA



David Wray, ACA, CPA, CGA, MBA, BFP, held finance executive roles in the technology industry for over 25 years, is the President of DW Group. He is a Board Member at IESBA, and the International CFO Alliance (ICFOA) and chairs its ESG & Technical Accounting working group. He regularly shares his views on sustainability, ethics & governance, data management, digitalization, and change management. He writes and speaks internationally about ethics, sustainability, and the power of technology to connect finance and sustainability. David is author of an Amazon best-selling book The Power of Potential: A Straightforward Method for Mastering Skills from Personal to Professional. All thoughts are my own, no AI created content here!

Anchoring Ourselves

We're living through incredibly strange times! Globally, traditional allies are feeling pressure to compromise their values and priorities for the benefit of others, tariffs are used as illegal pressure tactics, as affirmed by the majority decision from the US supreme court on February 20, 2026, and some country sovereignties are increasingly under threat. National governments are diluting environmental, social or governance requirements on companies – in some cases applying sanctions when businesses seek to remain responsible.

The result is a collision of politics, business and ideology, often with dramatic effects. As if this paradigm shift is not enough to manage, the battle for technology supremacy is also in full swing.

A growing tension is emerging between technology development, particularly for artificial intelligence (AI), intellectual property and data control rights and sovereignty. A recent example is Europe's decision to create its own independent technology solutions to protect the region's businesses from foreign political interference, amongst other things. An example cited by the EC was the risk of a foreign government directing one of its domestic businesses to disable technology usage access to its customers within a designated foreign country. Such a decision would have a severe and disruptive economic impact on the targeted country and its businesses. Such actions, previously unimaginable, are now genuine technology and going concern risks. Political business risk is occurring domestically as well – Anthropic faces an existential threat to its business following its decision to disallow usage of its AI platform for military purposes in the US.

Knowledge of the ins and outs of intelligent technologies is inconsistent across the ecosystem from businesses through to investors – the opportunity to educate and raise the bar remain untapped.

Technology challenges can be distilled down from this macro lens to the business itself. While much is written about the opportunities that AI presents, far less is said about the business risks and threats that arise with the growing use of AI.

A Starting Point for Reflection

Let's start with businesses

Consider the range of company responses to a technology use policy – some have no policy on the use of AI powered tools while others have structured and formal ones. The reason clear policies are important is to manage data leakage risks, protect against improper data usage (e.g.: intellectual property (IP) right infringement) and breaches of data privacy rights. Although companies with structured policies are arguably less exposed to some of these risks, they are not immune. Shadow AI usage outside of a company-controlled environment, an age-old issue, remains real and often results in company confidential data within a public domain. AI is used for finance activities that range from scenario analyses to report preparation (such as business cases and technical accounting position papers) through disclosure notes, among other things. Intelligent technology overreliance also poses a challenge, particularly when the user is unable to pre-form a “reasonable guesstimate” on the expected AI outcome. Or is unable to assess the technology’s inner-workings (commonly referred to the black box issue) or sense check its output prior to use in decision making.

Now audit and assurance firms

If we extend the technology challenges a little further, consider audit and assurance firm development and usage of AI. The ethical threats here are equally interesting, one example being self-review. Imagine any one of the audit firms developing an AI solution that detected a particular type of fraud or anomaly in accounts payable management. If the AI training was done with inadequate data in its development, resulting in a hidden deficiency, would the firm ever question the AI output given that it trained and developed the solution itself?

Other risks emerge as well; a notable one surrounds whether AI use should be disclosed so that the recipient of the output is aware and can formulate his/her own view on its reliability. This interesting debate has some saying AI is merely a tool, not dissimilar to using Excel, which, therefore, doesn’t require specific disclosure to an information user. Others take a view saying that, given intelligent technology risks, such as AI hallucinations, disclosure is important so users can determine for themselves the reliability and relevance of the information output.

Given intelligent technology risks, such as AI hallucinations, disclosure is important so users can determine for themselves the reliability and relevance of the information output.

A recent example in Australia: a large audit firm used AI to prepare a report for its government client, use of which was undisclosed in the report. The client subsequently detected errors within the report (i.e., including references and citations to non-existent reports by academics at the universities of Sydney and Lund in Sweden, aka hallucinations). This, coupled with the lack of transparency, resulted in the firm partially refunding its client on the circa 430K C\$ contract value. Seeing specific examples like this make it difficult to argue against greater transparency and more robust governance on the use of intelligent technologies within a business.

A fundamental question all consulting and advisory business entities should ask themselves is: *“why would a client pay top dollar for something generated by AI when the client could do near enough the same thing at a fraction of the cost?”*. The challenge will require human-centred insight to zero in on the value-added insights that are not attainable to its clients but for its human competencies and capabilities.



Completing this loop by introducing the future of work

Pushing this topic even further, consider how the use of intelligent technologies is shaping the future of work. Accounting and finance entry-level jobs are in dramatic decline. For example, the World Economic Forum (WEF) predicted that the accounting and audit professions would experience a circa 15-20% job decline. The accelerating change underway within our profession is arguably more dramatic and impactful than

anyone may have experienced in their careers to this point.

[WEF’s Future of Jobs Report 2025](#) noted that companies are prioritizing human skills (aka people skills or soft skills) as the most important attributes going forward. Namely, *“analytical thinking remains the most sought- after core skill..., with seven out of 10 companies considering it...essential...followed by resilience, flexibility and agility, along with leadership and social influence.”*

By the time we reach 2030, WEF predicts four possible scenarios depending on the speed of workforce readiness (e.g.: training) and the speed of AI advancements. These four outcomes (see Diagram 1) are:

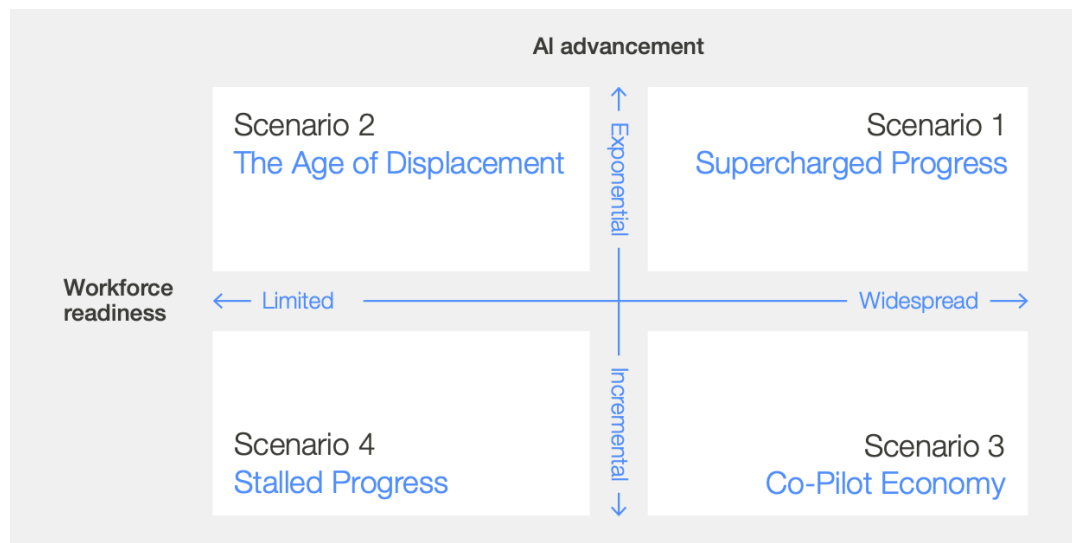
- Supercharged
- Displacement
- Co-piloted
- Stalled

Once we understand that AI is a tool designed to augment and perform routine tasks well, we can redirect our focus towards learning how to use it properly.

The significance of the WEF predicted models is the differing effects they will have on productivity, automation, labour market participation, inequality, consumer confidence, unemployment rates and operating margins, among other economic measures. The impact

on people and business is direct and will ultimately extend into societies, affecting things like healthcare services, retirement programs, cost of living and so much more. None of us can afford to be complacent in this transformation!

Diagram 1: WEF Four Scenarios for the Future of Jobs in 2030



Source: World Economic Forum.

The questions finance leaders should rapidly be asking themselves include, for instance, understanding how their team members can accelerate the acquisition of critical thinking and experiential insights that come, in large part, over time, with multiple and diverse hands-on experiences and by learning through mistakes or errors. How can new skills be accelerated to replace the approximately 40% of finance skills that will be in decline over the next five years? This connects with an essential ethical skill that is mind-set and the ability to stand back and consider wider ramifications, such as relevance, reliability, impact and consequences, before decisions are taken.

Or consider how accounting and finance functions may look in the future where intelligent tools, presently, create the misguided impression that the value of a finance leader (or any other functional leader, for that matter) is distilled into a process that can be automated. When we consider that a significant part of the value of finance leadership is rooted in ethics, innovation, creativity, human-centred complex analytical decision making and empathy, it is unrealistic to believe that AI will replace such leadership needs when AI lacks consciousness, self-awareness or an ethically centred understanding – all sentient human attributes. Sentient AI is an entirely different conversation given that it would radically upend life as we know it – something *homo sapiens* are clearly not ready for if we look to the current state of world affairs!

Work will evolve, something that we should not fear but rather embrace as an opportunity for a better future (just as other disruptive industrial changes have largely improved quality of life over time). Technology can do some things better – quickly perform queried analysis on very large data sets, prepare draft reports (based on pre-defined parameters or requirements) or any number of repetitive or predictable activities that can be distilled into a series of process steps that require little human intervention. Humans do some things better – apply judgment, understand complex issues, act ethically, create or innovate something new or act with empathy. Once we understand that AI is a tool designed to

augment and perform routine tasks well, we can redirect our focus towards learning how to use it properly – just as was previously done when Excel and ERPs replaced physical ledger book management.

Reality Sets In

Intelligent technologies are undeniably here to stay; they are an important element of the digital era. AI can significantly improve efficiency and efficacy within finance teams which, in turn, supports businesses and profitability growth – both important objectives for economic and social wellbeing.

To achieve the full potential that intelligent technologies offer, we need to start being more thoughtful and deliberate in ensuring that technologies are ethically developed and similarly used within organizations; and in better understanding the third-party software solutions that drive business and finance operations. To better understand the risks technology brings, in addition to its opportunities, well-rounded skills and capabilities will be necessary for the digital age leader.

To underpin individual and collective success, we must ensure that ethics sits at the heart of everything we do within the finance function. The four ethical threats we can never lose sight of when we develop and use intelligent technologies include: self-review, self-interest, advocacy and intimidation (see Table 1).

Table 1: Threat Definitions

Table 1: Threat definitions adapted from the IESBA Code of Ethics, 2025 Self-interest	The risk that an individual’s interests (including financial) will inappropriately influence judgment or behaviour.
Self-review	The threat that an individual will improperly evaluate the results of a previous judgment made, or an activity performed within its organization, and upon which they rely when deciding something today.
Advocacy	The risk that an individual will promote a particular position of another person to the point that the individual’s objectivity is compromised.
Intimidation	The threat that an individual is deterred from acting objectively because of actual or perceived pressures upon them.

In addition to ethical principles, appropriate governance structures and investment in developing teams to deepen their understanding and responsibility in the use of any technology, need to be in place. Governance structures should include organizational elements and functions from the Board of Directors through to first line finance roles. Every person within an organization has a role to play, even though governance failures often occur because of a single point of failure, which then ripples into other links in the chain (such as a lack of oversight around ensuring data confidentiality is applied throughout, and post-AI training). Governance system design is generally the primary reason for a failure.

Considering the aspect of learning and development, *Forbes* author Wendi Safstrom (March 2026) found that in companies: “nearly three-quarters of knowledge workers globally

now use AI at work, yet 60% say they have not received formal training to use it effectively.” Additionally, she found that “AI spending is projected to rise 44% in 2026—while training budgets are expected to grow just 5%, and average learning time is actually falling—from 47 to 40 hours per employee.”

This investment in technology, coupled with the small (or dwindling) investment in human skills, raises other interesting ethical questions. How will employees satisfy their professional competence and due care ethical responsibilities in the face of this corporate dynamic? Can we expect employees to recognize data output errors, omissions or gaps when they lack an understanding of how the technology works or what its inherent risks are? The questions quickly run from here once one starts to reflect on the real risks that sit alongside the opportunities.

While many of the underlying changes required need to ultimately occur starting with early childhood education, this article focuses on organizational considerations.

A Way Forward

Now we have a better grasp of the independent research findings that help scope the magnitude of the primary risks, opportunities and gaps. So, what approaches can we put in place to ensure that finance teams are ready? Meaning readiness to assume responsibility for intelligent technology outputs and their impact, to transition from any redundant skill to grasp new competency requirements; and, ultimately, to prepare for a future where the profession will look quite different within a mere three to five years.

Let’s talk about ethics with technology

The primary ethical attributes required to ensure responsible use of technologies include integrity, objectivity, professional competence & due care, and confidentiality (see Table 2).

Table 2: Ethical Attributes and Risk(s) Linkage

Attribute	Risk(s) that the attribute address
Integrity	Risk of relying on inaccurate outputs: <ul style="list-style-type: none"> For example: fallibility, hallucinations; third-party AI tool reliance
Objectivity	Bias from data or incentives: <ul style="list-style-type: none"> For example: data dependence, scalability; overreliance on outputs; third-party reliance
Professional Competence & Due Care	Need to understand complex systems (opacity, non-determinism): <ul style="list-style-type: none"> For instance: qualified personnel approve outputs/assumptions, limitations and contrary information is documented
Confidentiality	Risks from interconnected systems <ul style="list-style-type: none"> Such as including sensitive data without permission, rely on outputs without permission or share unverified outputs externally
Professional Behaviour	Responsible use of technology

The first step in the responsible and ethical development and use of any technology is awareness of its risks, weaknesses and limitations. This approach links to the long-standing concept of “what could go wrong?,” regularly used by auditors and senior finance leaders when assessing risk environments. Adopting this mind set walks us through the myriads of questions that require answers for each phase of the technology development and/or usage. Some helpful questions to get started are:

- What could go wrong with our use of X intelligent technology (e.g.: AI)?
- What would AI failure look like?
- What are most significant risks we face in implementing this technology?
- What key assumptions have been made; what if they are proven wrong?
- What early warning signs would we see to flag us about an issue?
- What might cause us to be the authors of our own failure?
- What are the potential blind spots in our AI approach?
- What unintended consequences might arise from our AI project?
- What impact will the AI have on the organization (including customers, suppliers and partners), how are we ensuring readiness for that change?
- What legal or regulatory changes are being considered that could affect the project?

Additionally, the [Alan Turing Institute](#) provides a helpful “8 Questions You Should Ask First” checklist when considering artificial intelligence, and the EU has created a user-friendly website to help developers and users ensure their respective AI activity is compliant with the [European Union Artificial Intelligence Act](#) (aka EU AI Act). The EU AI Act is currently considered the gold standard in this domain.

Intelligent technology failures, when they occur, often harm customers, users and the reputation of the organization at the heart of the “scandal.”

Let’s talk about upskilling teams

Two considerations are relevant when reflecting on skills – are they academic skills in preparation for a career role (most often delivered through university or other educational organizations) or professional level skills for those already in the profession and looking to pivot or progress (most often delivered through professional accounting organizations). CPA Canada actively advocates for AI guardrails. In late 2025, CPA Canada President and CEO, Pamela Steer, stated in the letter to the Minister of Artificial Intelligence and Digital Innovation, Evan Solomon: “Ensuring that humans remain central to critical decision-making mitigates risk and fosters trust in AI outputs.”

In effect, AI training needs to purposefully elevate human oversight into the forefront of learners so that it becomes second nature. This is particularly important in areas where usage requires significant and complex ethical judgments, a high degree of accuracy and completeness, such as in financial services, accounting, assurance and legal professions among others.

So, we know it is important but are learning institutions doing something about it to embed responsible AI into its students? In a word, it depends. Many programs available, particularly those that are “free,” focus on using it right away without much, if any, weight given to ethics, critical thinking or the importance of accountability. Professional level programs in

the finance and accounting profession are, however, starting to actively embed and connect these concepts more directly in their learning (both certifications and PD courses).

Business leaders and individuals seeking to upskill their teams or themselves would be wise to reflect carefully on their professional responsibilities and ensure that the course or program they wish to select addresses the connectivity between technology and ethics, responsibility and accountability.

Bringing It Home

Intelligent technology failures, when they occur, often harm customers, users and the reputation of the organization at the heart of the “scandal.” We’ve seen dozens of examples of what can go wrong – from the [airline chatbot case](#) that ended up in court with a finding that held the airline accountable for its chatbot’s erroneous guidance to its customer or, as noted earlier, the [Accounting firm's customer report case](#) in Australia when the firm was required to partially refund its government customer for the hallucinations in its report. The PR fallout can be costly and have an impact on an organization well beyond settlement of the original intelligent technology failure liability.

The outcome would have been quite different if a robust assessment of the “what could go wrong” mindset had been used from the beginning, underpinned by an ethical mindset approach laying at the heart of its intelligent technology deployment.

It is critical that, as a profession, we lead in being part of ethical technology solutions as our future may well depend on it!



Standardization Meets Innovation: AI Platform for Regulatory Data

By Björn Fastabend



Björn Fastabend is head of the XBRL collection and processing unit at BaFin, Germany's Federal Financial Supervisory Authority, where he supervises all related activities and leads the implementation of strategic initiatives. He is also a member of XBRL International's Board of Directors, and previously chaired their Best Practices Board. The views and opinions expressed in this publication are those of the author. They do not purport to reflect the opinions or views of BaFin.

The Urgency Paradox

It is 2026 and the use of AI has taken hold in many industries. The private sector is leading the charge, adopting AI solutions to improve processes, cut costs and increase efficiency. Regulators are much more hesitant, since AI results are not yet deterministic and fully explainable — a significant concern when regulatory rigor is non-negotiable. This imbalance is not sustainable, however. Regulators simply do not have the luxury to wait for a perfect AI solution and must find ways to integrate AI into existing processes now.

The accessibility of AI has produced a flood of software solutions, each promising to solve a specific problem. It is tempting to implement one tool for summarizing PDFs, another for analyzing XBRL filings and yet another for scanning social media for early warning signals of market instability. The issue with this approach is straightforward: rushing into AI adoption without a coherent strategy easily leads to a "zoo" of disconnected applications, each with its own requirements, maintenance overhead and integration challenges — technical debt that will need to be managed for years to come.

So how can this be avoided? This article presents a thesis that is both straightforward and, at first glance, counterintuitive: standardization. Not as a barrier to innovation, but as the very foundation upon which sustainable innovation can flourish.

View standardization not as a barrier to innovation, but as the very foundation upon which sustainable innovation can flourish.

Beyond Format: The Database Foundation

Regulators commonly receive large quantities of reports from the companies they supervise. Depending on the size of the country, this can mean thousands of files for a single reporting deadline and a single mandate. In order to manage and process all of these, structured reports are disseminated into a database — the foundation upon which BI tools analyze the data.

I have seen many discussions about which format is best suited for AI analysis: XBRL, XML, JSON, CSV, even Excel. But these debates miss the point entirely. It is extremely rare that regulators analyze individual files directly. This would be far too inefficient given the volumes involved.

The real question is: which format preserves the most semantic richness on its journey to the database, giving an AI the best possible baseline for thorough analysis? XBRL is the gold standard here. Its taxonomy encodes meaning — not just structure — and can be stored in the database alongside the reported data. That said, once any report has been disseminated into a database, it is implicitly in a structured format, since the database schema provides structure of its own, regardless of the original file format.

This leads to the first and most fundamental action point: a unified database access layer. By standardizing how data is accessed — regardless of which schema it lives in — regulators eliminate the need to reinvent the wheel each time a new reporting format or mandate is introduced.



The Platform Imperative

Let's revisit our "AI zoo" scenario. A regulator, caught up in the rush to adopt AI, has ended up with 10 individual applications: one for sentiment analysis of social media, one for XBRL analysis, one for PDF summarization, and so on. Each is specialized, each is valuable — but together they represent a significant maintenance and interoperability burden. And, while each application has its own specific use case, they all draw from the same underlying data sources.

This is precisely where platform thinking changes the equation: build common infrastructure once and use it many times.

Without getting into technical details, a well-designed AI platform for regulatory use would standardize four core components:

Database Access ensures that all AI applications connect to the various databases through a consistent, standardized interface. The structure and content of each database may differ, but the way applications access them should not.

Model Management handles the AI models themselves. Different use cases will require different models — analyzing structured XBRL data calls for a more specialized, fine-tuned model than summarizing a PDF. But there is no reason for each application having its own approach to hosting and managing those models. Standardizing this infrastructure means the platform is ready for both current and future applications.

Governance & Audit is non-negotiable in a regulatory context. Every analytical step must be traceable: who queried what, when and with what result. User access controls, version tracking and prompt and response logs are universal requirements for any AI application used in regulatory supervision — and therefore natural candidates for centralization.

Quality & Validation addresses the elephant in the room: AI hallucinations. Regulators simply cannot afford to act on inaccurate results. One practical approach is to deploy a second model to fact-check the first — only surfacing results to the analyst when a confidence threshold is met. While the specific validation method will vary by application, the interface for reporting confidence scores and flagging uncertain results can be standardized across the platform.

The overall analogy is a paved road versus everyone building their own path. The initial investment is greater, but the long-term benefits in maintainability, extensibility and reliability far outweigh the cost of building each individual path from scratch.

The Payoff: Innovation Without Chaos

So, what does this actually look like in practice? Let's walk through a concrete scenario. A regulator already has XBRL processing up and running on its AI platform — meaning the four standardized components are in place: a unified database access layer, model management, governance and audit infrastructure, and quality validation.

Now a new CSV-based reporting mandate comes in. Without a platform, this would mean building a new standalone application from scratch: designing database connectivity, setting up a new model, implementing governance controls and building a validation mechanism. With the platform already in place, the picture looks very different. A new database is set up to store the CSV data, and a specialized model is selected for analysis. But the new model slots straight into the existing model management system. The database access layer handles connectivity without custom integration work. The governance and audit infrastructure is already there. And the validation mechanism works in a very similar way to XBRL processing, since both are dealing with structured data. What would previously have been a significant development effort is now a targeted extension of existing infrastructure.

The same logic applies when a better model becomes available for an existing use case — simply swap it into the model management system and the rest of the application continues unchanged. Likewise, if a vendor change becomes necessary, the platform's standardized interfaces mean individual components can be replaced without disrupting existing functionality.

The pattern that emerges is clear: the first use case requires real investment, but each subsequent one benefits from everything that came before, making it progressively faster and cheaper to expand. The key insight is simple — let innovation happen in models and applications, not in infrastructure. Good

infrastructure should be nearly invisible, quietly enabling everything built on top of it. That is the true power of standardization!

The beauty of standardization is that it's a strong, mature foundation upon which true innovation can take place.


Practical Path Forward

Right now, almost everyone wants to get their hands on an AI application. This increases the risk of inadvertently creating an "AI zoo" that will be much more difficult to maintain, expand and innovate in the future. While it is likely overkill to jump right in and design an AI platform when it is not even clear if there will be multiple use cases for AI applications, as soon as there are several potential AI apps on the horizon, it's not too late to plan bigger and with a standardized platform in mind.

That said, even when building the first AI app, it already makes sense to plan for scalability from day one. It doesn't yet have to be the full platform, but how data is being fed to the AI is already a significant decision that will potentially have an impact on future expansion. Instead of feeding individual files directly to the model, invest in storing the data in a database. Introduce common access patterns for various components, making it easier to extend future parts.

The key is to let innovation happen in models and features, not the underlying infrastructure. The beauty of standardization is that it's a strong, mature foundation upon which true innovation can take place. Good standardization is nearly invisible, while innovation is what is seen and noted. Think about shipping containers. Nobody ever thinks about how a global marketplace that enabled concepts like "just in time delivery" for manufacturing came to be. Yet, the introduction of the shipping container had a significant impact on this. It's the stable baseline that provides the foundation for innovative concepts, something that was simply too expensive prior to standardized shipping.

Avoid the "it worked once but doesn't scale" trap that is so easy to fall into when first introducing a new technology.



Why This Matters

When considering the introduction of AI in any workplace, but especially that of regulators, it is essential to think carefully before making that move. Consider all possible use cases up front, even those that may only materialize further down the road. Plan for the long run, not the short wins. The introduction of AI will have a disruptive impact even on regulatory work, and planning sustainable solutions instead of one-off experiments is the key.

We must also consider that we have been truly exposed to the power of LLMs for only a very short time. ChatGPT was made available to the public in 2022 and now, just three years later, the pace of development gives a clear glimpse of what might come. Without a solid foundation, it will be much more difficult to keep up. Introducing and benefiting from new features and developments in AI will be significantly harder if there is no stable infrastructure that allows for the replacement of parts or upgrades of models.

Furthermore, as regulators, we have a unique advantage that private sector competitors do not: we don't compete with each other. We can learn from each other's experiences rather than reinventing the wheel numerous times. International forums and collaborative initiatives are invaluable in this regard.

While I am a firm believer in the power of standardization, I frequently hear skepticism that it stifles innovation. Think about shipping containers. Nobody ever thinks about how a global marketplace that enabled concepts such as "just in time delivery" for manufacturing came to be. Yet, again, the introduction of the shipping container had a significant impact on this. Or take internet protocols such as TCP/IP, which enabled the world wide web. When was the last time you wasted a single thought on TCP/IP while in a video conference? In both cases, it's just there and it just works. That is the true beauty of standardization done right — when the infrastructure that supports great innovative solutions blurs quietly into the background.

So let your innovative solutions take center stage, while the platforms and infrastructure that enable them to take a seat in the back, providing the stability and maturity needed to let innovation shine.



Beyond Compliance: How AI and XBRL Are Reshaping Financial Analysis

By Janis Steinmann



Janis Steinmann is Head of XBRL at Lucanet, leading the company's XBRL team while positioning XBRL solutions as a cornerstone of Lucanet's CFO platform. He also serves as Chair of XBRL International's Best Practice Board and member of ESMA's Disclosure Working Group within the Data Standing Committee.

Annual reports published once a year won't satisfy investors using AI-powered analysis tools that expect continuous, structured data feeds. For two decades, XBRL has provided the infrastructure to make machine-readable financial data possible, standardizing reporting across 70 countries and covering tens of millions of companies.¹ Most companies, though, have treated it as a compliance checkbox: file the report, move on.

That approach no longer makes sense. Large language models can now query XBRL-powered databases in natural language, transforming how finance teams analyze data. At Lucanet, we've built an XBRL Tagger Agent that automates the creation of structured reports, including narrative disclosures. This shift affects both how companies create XBRL data and how analysts and investors consume it.

This article examines what happens when AI meets structured financial data: The practical challenges, the current limitations and why data quality has become a competitive issue rather than just a technical concern.

Large language models can now query XBRL-powered databases in natural language, transforming how finance teams analyze data.

The Challenge: Why Financial Data and AI Need Each Other

Financial reporting requires deterministic answers. Ask "What is our turnover for the financial year?" and you need one correct number. AI systems work differently: They are probabilistic. Ask a large language model the same question five times, and you might get five different narratives around the same core data. In regulated finance, this creates problems.

The real challenge is context: A database column labeled "period" with integer values 1-12 means nothing without additional information. Is it a month? A fiscal period? Without a semantic layer explaining what the data represents, AI systems may send incorrect queries to databases or misinterpret the results.

Modern annual reports run 200-300 pages, with some exceptions going up to 1,000 pages. AI systems have processing limits: Think of it as desk space. You can fit 50 pages on your desk. Someone asks about page 51, and you have to clear everything else, losing context. Upload a full PDF to an AI system, and it either truncates the document or loses critical information buried in the middle sections.



PDFs compound this problem further. They present visual formatting, not structured data. The AI must infer what numbers represent, which time periods they cover and how concepts relate to each other – and it may infer differently each time. Ask the same AI system the same question about a PDF document tomorrow and you may get a different answer. The probabilistic nature of language models means that repeated extractions from unstructured sources produce varying results. The effect is markets running on probabilistic interpretations rather than deterministic facts.

XBRL-based systems can solve this through selective retrieval. Instead of forcing an entire document through a narrow processing window, the system queries specific data points as needed. Each fact carries dimensional context: entity, time period, measurement unit, and additional qualifiers. The AI doesn't need to hold everything in context. It asks for what it needs, when it needs it.

XBRL provides machine-readable data with embedded definitions. Each concept comes with authoritative metadata, for example the thousands of concepts in the US GAAP taxonomy alone, each properly defined.² The taxonomy structure explains relationships between concepts. Time dimensions specify reporting periods. Ideally the company's management signs off on this structured data, making it a single authoritative source.

This infrastructure matters for AI because it eliminates ambiguity. When an AI system queries "Revenue" from an XBRL-based database, it gets a precise definition, the calculation methodology, the reporting period and the dimensional context. No guessing required.

AI in XBRL Creation: Automating the Tagging Process

Companies spend thousands of hours preparing annual reports: gathering data, consolidating numbers, reviewing drafts, securing sign-offs, arranging assurance, polishing design. The XBRL tagging? Typically, the last in this workflow, it requires a lot of manual tagging labour, thus many finance teams treat this final step as burdensome compliance work. File it and forget it. The disconnect is striking. Companies invest heavily in the report itself but minimize attention to its digital representation. This made sense when XBRL served primarily regulatory purposes. It makes less sense now as the reports become more accessible and AI systems will start to use that structured data as an analytical source.

We've automated most of the tagging process through the XBRL Tagger Agent. It handles taxonomy mapping automatically, including narrative disclosures, not just the numbers. LLMs can match report content to taxonomy concepts with reasonable accuracy, in a fraction of the time manual tagging requires. That said, using prompts and a model alone will not produce reliable results across a full report. Consistent automation is achievable only because of publicly available, structured financial data in XBRL, which the LLM uses to identify appropriate tags, not only in tables, but in the notes section as well.

The practical reality still requires a human-in-the-loop. The AI completes the bulk of the work: identifying financial statement line items, mapping them to appropriate taxonomy elements, handling dimensional breakdowns. The human is still responsible for the data that will be published, similar to using a Tagging as a Service agency, the work will have to be reviewed and signed off on.

We've also experimented with using LLMs to consume XBRL data, testing how well these models can query structured repositories and generate analytical insights. The results suggest that both creation and consumption of XBRL data will increasingly involve AI systems.


When your structured data feeds AI-powered analysis across the investment community, mistakes become competitive disadvantages.

Why Data Quality Matters More Than Ever

Phil Wilding, who helps portfolio companies at HG Capital implement AI, put it directly in a recent *Finance in Flow* podcast episode: "AI doesn't fix bad data, it amplifies it at scale." A tagging error in your XBRL file gets picked up by dozens of AI analysis tools, propagated through comparative databases and incorporated into investment models. Your structured data becomes a broadcast signal, and errors travel as fast as accurate information.

This changes the stakes. When XBRL was primarily a regulatory filing, mistakes were embarrassing but contained. When your structured data feeds AI-powered analysis across the investment community, mistakes become competitive disadvantages. Your company's beacon broadcasts static instead of signal.

The solution requires better governance. Independent assurance of digital reports should become standard. AI-assisted preparation and review tools can help: they can flag inconsistencies, check calculations, benchmark against peers. But companies need to treat their digital disclosures with the same rigor they apply to the annual report itself.



AI in XBRL Consumption: Conversational Analysis

Model Context Protocol (MCP) servers enable AI systems to query XBRL-based databases through natural language. Instead of learning technical interfaces, analysts ask questions: "Compare liquidity ratios over the past year." The system translates this into appropriate database queries, retrieves structured data, performs calculations and presents results.

Think of MCP servers as specialized assistants sitting between the AI and various databases. When you ask an AI to analyze a company's financial ratios, it recognizes this requires external data and contacts the XBRL-focused MCP server. The server accesses a financial database, locates relevant company information and returns results with appropriate context. The AI combines this with its knowledge of financial analysis to provide a comprehensive response.

This capability depends on accessible repositories of XBRL data. The SEC's EDGAR system maintains structured financial filings from US public companies going back more than a decade: a petabyte-scale repository of authoritative financial information.³ Similar repositories exist in Japan, China, India, Korea and across Latin America. Each represents a rich holding of public information that AI systems can interrogate.


Europe is working toward an integrated approach through the European Single Access Point (ESAP).⁵ This mechanism will consolidate access to financial and sustainability reporting across EU member states, creating a unified repository for structured business data. As these repositories expand globally, they provide the data infrastructure that makes conversational financial analysis practical.

MCP servers address this directly. The complete XBRL dataset stays in an external database rather than being forced through a processing window. When you ask "Show me the profit and loss for Hays," the system extracts precisely the relevant facts, performs calculations and returns a focused response. You can compare metrics across dozens of companies, analyze trends over multiple years and perform sophisticated modeling, all within a single analytical session.

Every piece of information carries precise metadata about its origin. When the system calculates a debt-to-equity ratio, it specifies exactly which balance sheet line items were used, from which reporting period, and references the specific XBRL tags that define those values. This granular attribution extends to complex calculations and comparative analyses.

José Antonio Huizar Moreno, a member of XBRL International's Best Practice Board, developed an MCP tutorial demonstrating how such systems can trace any calculation back to its source. Click on a revenue figure and you can see exactly which company's XBRL filing it came from, which calculation was applied and that the values are taken directly from official filings, not inferred by the AI.

This traceability matters for compliance and audit. Analytical work can be documented, reviewed, and validated efficiently. Rather than manually searching through PDF documents to verify AI-generated insights, analysts can immediately trace the logical path from source data through calculations to conclusions.



Trust but Verify

John Turner, CEO of XBRL International, uses the phrase "trust but verify." Use AI to discover new relationships and spot anomalies across datasets. These tasks would be difficult manually. But check that the AI answers are correct using conventional analytic tools. AI isn't magic. The systems are designed to be imaginative, which doesn't align well with financial analysis requirements.

Human oversight remains essential for critical decisions. Implementing AI-powered financial analysis requires explainable systems from day one: clear audit trails showing LLM outputs, inputs, reasoning, logic and source data. Can a human reconstruct the AI's path? If not,

don't put it into production.

Think of AI as a junior analyst on your team. Give it the right training, provide context about the company, supply the numbers and the task. Someone senior validates the work and checks it. If it's wrong, you discuss it, refine the output. You don't let a junior analyst make critical business decisions unsupervised. The same principle applies to AI systems.

Current Challenges and the Competitive Reality

The most significant limitation is timing. Annual reports are published once per year. Investors and analysts need ongoing information. The SEC requires quarterly XBRL filings for US public companies,⁴ but interim structured data remains limited in many regions. Europe's ESAP will

help standardize access across member states, but the fundamental mismatch remains between AI systems expecting continuous data feeds and reporting frameworks built around annual cycles.

This creates challenges for cross-border analysis as well. While repositories like EDGAR provide comprehensive coverage of US markets, analysts comparing companies across jurisdictions must navigate different filing frequencies, varying taxonomy standards and inconsistent availability of structured data. ESAP should help address this for European markets, but global analysis still requires working with multiple systems and standards.

Data quality suffers from an "everyone and nobody" responsibility problem. The CFO signs off but may not understand the digital layer in detail. Auditors check human-readable reports but often skip the structured data entirely. Finance teams think it's an IT problem. IT teams think it's a finance problem. What's needed is someone who treats every XBRL tag as if it will be analyzed by a thousand algorithms tomorrow. Because it will be.

Independent assurance should become standard. Auditors need to verify structured data, not just documents. They should confirm that when something is tagged as revenue, it actually represents revenue and not some other concept entirely. Without this verification layer, the quality problems compound.

Some companies will implement AI-assisted workflow improvements incrementally, finding efficiencies in data collection, analysis, validation and reporting. These organizations become more competitive.

The Competitive Reality

John Turner's assessment is direct: "We're going to witness a transfer of wealth from companies that treat data as paperwork to those that treat it as power."

The mechanism is straightforward. Some companies will implement AI-assisted workflow improvements incrementally, finding efficiencies in data collection, analysis, validation and reporting. These organizations become more competitive. Others will continue treating digital reporting as a compliance exercise. They fall behind.

The media focuses on AI's impact on jobs. Turner sees a different dynamic: "I'm not convinced there will be a substantial impact in that area. People are adaptable." The real impact hits at the corporate level. Creative destruction accelerates. Winners and losers emerge faster.

Within five years, AI systems will have worked through virtually every piece of structured financial data that is publicly available. Companies filing poor-quality digital reports risk becoming invisible to these systems. Data quality, in that sense, becomes a form of institutional reputation, and that shift is already underway.

Conclusion

XBRL isn't outdated. It is essential infrastructure for the AI age. John Turner's analogy in a *Finance in Flow* podcast episode¹ captures it: digital reporting is plumbing – invisible, but

critically important. The combination of structured data and AI enables analysis that was previously impractical at scale.

Financial reporting will increasingly look less like documents and more like data services: verifiable, continuously updated, with enhanced assurance built in. Companies will spend less time reconciling numbers as the underlying infrastructure becomes more trustworthy, and more time on the interpretive work that actually requires human judgment. Internal transformation will happen faster than external regulatory change: internal analysis and benchmarking will shift rapidly as structured external data becomes easier to consume automatically.

The technology to do this exists now. AI-assisted tagging is in production. MCP servers that connect structured XBRL repositories to conversational analysis tools are straightforward to deploy. What remains is largely organizational: whether companies will treat their structured data with the same seriousness as the document it accompanies. Turner's framing is hard to argue with. The transfer of wealth from companies that treat data as paperwork to those that treat it as power is not a prediction, it is already underway.

Sources

¹ John Turner, CEO, XBRL International. *Finance in Flow* podcast, 2024. <https://youtu.be/Jp6E38TlaZk>.

² FASB GAAP Financial Reporting Taxonomy. Financial Accounting Standards Board, 2025. <https://fasb.org/xbml>.

³ About EDGAR. U.S. Securities and Exchange Commission. <https://www.sec.gov/submit-filings/about-edgar>.

⁴ Inline XBRL. U.S. Securities and Exchange Commission. <https://www.sec.gov/data-research/structured-data/inline-xml>.

⁵ Regulation (EU) 2023/2859 of the European Parliament and of the Council of 13 December 2023 establishing a European single access point. Official Journal of the EU, 20 December 2023. <https://eur-lex.europa.eu/eli/reg/2023/2859/oj>.



The Fusion of Risk Management for Artificial Intelligence and Human Agents

By Eric E. Cohen, CPA



Eric E Cohen, CPA, is a technologist with a passion for collaboration toward the goal that “a piece of business information, once entered into any system, anywhere, never needs to be retyped as it moved through the business reporting supply chain.” He’s also a prolific author, engaged in virtually every effort to standardize accounting and audit data, a national expert to a wide variety of standards efforts, and co-founder of XBRL.

The Fall 2023 *ThinkTwenty20* included an article I wrote called “Generative Artificial Intelligence and Undesirable Output: Developing a Taxonomy for Undesirable Output with a Focus on Hallucination.” It spoke of my ongoing and personal campaign to move beyond the word “hallucination” when blaming GenAI for the problems of undesirable results. “Hallucinations” is such an easy term, but it is an oversimplification, and it masks the real problems, which are often not a quality issue with the AI at all.

Consider a period-end close process in 2026. A controller issues a natural language instruction to an AI agent, which delegates sub-tasks to two additional agents, one reconciling accounts receivable, the other performing foreign currency translation. The agents execute journal entries, flag exceptions and deliver a draft close package. The process is faster, more consistent and impressively automated. Now ask a simple governance question: who is accountable for the result?

If the answer requires reconstructing a chain of events across an AI agent log, a human identity system, an ERP audit trail and a task management platform (none of which share a common schema), the governance framework has a fundamental gap. That gap is the subject of this article.

This article describes a collaborative effort to create something called the *Agent Metadata Specification*. The work is an open-source project, with a goal of supporting consistent and more formalized agent risk management, promoted by an organization called Tavro AI, LLC (Tavro) and involving numerous contributors from industries such as banking, healthcare and

energy and utilities.¹ A GitHub repository² is the shared home for the technical deliverables. The author of this article expanded on that work³ with a critical addition: a parallel treatment of human actors and the interactions between humans and AI agents that the original specification left unaddressed. The goals of the original work and the author's expansion are the body of this article.

What Is the Need for an Agent Metadata Specification?

The Agent Metadata Specification was created to address the need for formal risk management tools and methodologies in light of the proliferation of artificial intelligence agents. While agents in computing are not new, agents supported and driven by large-language model (LLM) based generative artificial intelligence tools have brought powerful and potentially risky autonomous capabilities to individuals and organizations at a scale and speed that governance frameworks have not kept pace with.

Agents supported and driven by large-language model (LLM) based generative artificial intelligence tools have brought powerful and potentially risky autonomous capabilities to individuals and organizations.

Tavro characterizes the result as an "AI Agent Visibility Vacuum," driven by a phenomenon they term *shadow action*: the unmanaged and autonomous multiplication of AI agents across enterprises, mirroring the earlier governance challenge of shadow IT. The consequences are concrete: organizations do not know that agents exist, who owns them or where they operate; agents often function as unmonitored "digital insiders" with high privileges and broad access to sensitive data and systems; and redundant, unmapped agents waste compute cycles and inflate cloud costs.

For enterprises operating within the European Union, the stakes are further elevated. This visibility vacuum makes it extremely difficult to conduct the consistent, evidence-based risk assessments mandated by the EU AI Act, and to demonstrate continuous compliance. Current agent platforms lack a unified, standardized mechanism for exporting the governance data that regulators require.

Are AI agents New?

Those who have been engaged in AI in accounting information systems for many years will point to today's agentic evolution as only the latest in a long line of agents in AI. One key correspondent is Miklos Vasarhelyi; he and I discussed how new and novel the concept of intelligent AI agents in accounting actually is.

¹<https://www.tavro.ai/agent-metadata-specification-is-an-open-source-project-to-support-consistent-agent-risk-management/>.

²<https://github.com/TavroOrg/Agent-Metadata-Specification>.

³<https://github.com/GenAICPA/Agent-Metadata-Specification>.

For those unfamiliar with him, my colleague Miklos is the KPMG Distinguished Professor of Accounting Information Systems and Director of the Rutgers Accounting Research Center and Continuous Auditing and Reporting Lab (CAR Lab). To give some perspective on his familiarity with AI agents: the 70th World Continuous Auditing and Reporting Symposium (WCARS) is scheduled for May 21–22, 2026, in Seville, Spain – a conference I am scheduled to address on the topic "From Generative AI to Agentic Audit: The Rapid Evolution of Tooling for Continuous Auditing and Reporting." The first conference of the series, held in 1998, was already focused on online audit and included a discussion on intelligent agents.⁴ As reflected in his curriculum vitae,⁵ Miklos has written, spoken and taught about intelligent agents in accounting numerous times, and EY hosted a conference on "Agents and Online Audit" as far back as 1997.

Miklos told me that, in his view, what we call "agents" today are a natural evolution of what has been part of computing almost from its inception – pointing to daemon (an automated background process in operating systems) and cron (a time-based job scheduler), autonomous helpers that have long managed system resources and tasks on behalf of users and systems.

Other founding contributors to continuous audit have also championed agents. The late Michael Groomer and Uday Murthy wrote about the role of agents in embedded audit modules in their paper "A Continuous Auditing Web Services Model for XML-Based Accounting Systems"; I was the discussant on that paper when it was presented at the University of Waterloo Centre for Information Integrity and Information Systems Assurance (UWCISA) conference some 20 years ago.⁶ Likewise, Roger Debreceeny, along with his long-time collaborator Glen Gray and their co-authors, championed agents in work such as "The Development of Embedded Audit Modules to Support Continuous Monitoring in the Electronic Commerce Environment."⁷

If AI Agents Aren't New, What Has Changed?

While Miklos and I may have a friendly disagreement about how different today's agents truly are from those of the past, there is a wide range of both technical and practical differences between agents in 1997 and those in 2026.

Consider the practical issues first. What were the challenges of embedded audit modules leveraging software agents 30 years ago?

⁴Program available through the Wayback Machine at the Internet Archive, at https://web.archive.org/web/20220618034342/http://raw.rutgers.edu/docs/wcars/1wcars/1WCARS_Agenda.pdf.

⁵<https://www.business.rutgers.edu/sites/default/files/documents/faculty/cv-miklos-vasarhelyi.pdf>.

⁶<https://www.sciencedirect.com/science/article/abs/pii/S1467089504000156>

⁷Roger Debreceeny, Glen L. Gray, Wai-Lum Tham, Kay-Yiong Goh, and Puay-Ling Tang. "The Development of Embedded Audit Modules to Support Continuous Monitoring in the Electronic Commerce Environment." *International Journal of Auditing*, 2003, 7, pp. 169–85.

- **Integration barriers.** Agents were not integrated into the fabric of applications. Introducing an embedded audit module meant that auditor-specific software would have to be installed on a client's systems, requiring bespoke implementations at each engagement and creating unhealthy interdependence between client and auditor. In 2026, agentic tools are embedded throughout commercial platforms; SAP's Joule, for example, not only integrates AI across the solution but provides end-user, citizen-developer-level agent building features.
- **Purely programmatic design.** To create an agent, you needed developers. Today's agents can be placed into service by non-developers using wizard-driven tooling and natural language interfaces.
- **Limited capability.** Even the robotic process automation (RPA) tools that emerged in the 2000s and became popular in the 2010s were limited to time- or activity-triggered steps following prescribed rules. A 2026 AI agent is qualitatively different. It can take natural language instruction on goals to be accomplished; evaluate the tasks required and develop plans; consider available tools; seek clarification where necessary; share plans and proceed when acknowledged; use tools and frontier model knowledge to perform tasks and evaluate progress – with tools including screen interaction, application programming interfaces (APIs), guidance through the Model Context Protocol (MCP), and orchestration of other agents; provide progress feedback as requested; and create and deliver required outputs.
- **Infrastructure immaturity.** Accounting information systems and the broader IT environment were simply not ready to support sophisticated agents at scale. I have long cited the work of Professor George H. Sorter, who proposed "Events Theory"⁸ more than 50 years ago – the idea that enterprises should expose business events at a detailed level and let stakeholders interpret, aggregate and use the information as their decision models require, rather than as management anticipates. Barriers such as storage costs, communications infrastructure, data standardization and AI analytical capability have steadily fallen, making the event-level exposure Sorter envisioned increasingly achievable.

Have the Changes Brought Significant Activity by Agents?

While solutions to build agents existed before the generative AI era, giving non-developers meaningful control over computers through frontier-level AI is relatively new. My own experimentation began with Anthropic's public beta of Claude for Computer Use⁹ using Claude 3.5, in October 2024. This launched the ability for non-developers to control a computer desktop through a graphical user interface – the agent would take screenshots to "see" the screen, then move cursors, click buttons and type text to complete multi-step tasks. In just 18

⁸<https://home.business.utah.edu/u0030831/Sorter%20AR%201969.pdf>.

⁹<https://www.anthropic.com/news/3-5-models-and-computer-use>.

months, we have moved from that to Claude Cowork,¹⁰ which advertises the capability to perform a complete period-end close, involving a human only when judgment is required.

Google released Gemini CLI¹¹ in June 2025 and previewed their own Computer Use model in October 2025.¹² OpenAI released ChatGPT Agent in July 2025 with its own virtual browser environment. Microsoft partnered with Anthropic to release Microsoft Copilot Cowork,¹³ which Microsoft describes as possessing business understanding including procedural knowledge from business workflows and the ability to work across structured and unstructured accounting data from Dynamics 365 and other sources.¹⁴

Analysts confirm that agents are rapidly becoming widespread. As of November 2025, McKinsey reported that 60% of survey respondents were at least experimenting with agents.¹⁵ PwC reported a higher figure – 79% of respondents had already adopted AI agents, with 17% reporting full adoption throughout the company.¹⁶ The Cloud Security Alliance (CSA) reports that 43% of organizations say more than half of their employees use AI agents regularly.¹⁷

The same CSA survey, however, reveals the governance reality behind those adoption numbers. Between one and 100 shadow AI agents – unsanctioned deployments where ownership is often unclear – have been reported by 54% of organizations, and 44% report low or no confidence in

¹⁰<https://www.anthropic.com/product/claude-cowork>.

¹¹<https://blog.google/innovation-and-ai/technology/developers-tools/introducing-gemini-cli-open-source-ai-agent/>.

¹²<https://blog.google/innovation-and-ai/models-and-research/google-deepmind/gemini-computer-use-model/>.

¹³<https://www.microsoft.com/en-us/microsoft-365/blog/2026/03/09/copilot-cowork-a-new-way-of-getting-work-done/>;
<https://techcommunity.microsoft.com/blog/microsoft365copilotblog/a-closer-look-at-work-iq/4499789>.

¹⁴<https://www.microsoft.com/en-us/dynamics-365/blog/business-leader/2026/03/09/a-new-way-of-working-is-taking-shape-frontier-transformation/>.

¹⁵<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.

¹⁶<https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html>.

¹⁷<https://www.businesswire.com/news/home/20260416255682/en/More-Than-Half-of-Organizations-Experience-AI-Agent-Scope-Violations-Cloud-Security-Alliance-Study-Finds>.

their ability to detect AI agent-specific threats. CSA recently published their survey report, *Enterprise AI Security Starts with AI Agents*.¹⁸

How Are Agents Being Managed?

The candid answer is: generally, they aren't. Agents have come to the forefront of organizational work not through careful planning but through users finding tools to make their jobs easier, or managers seeking competitive advantage without IT involvement or management oversight. Because agents are now easy to build and deploy, they are more difficult than ever to discover. They operate with the rights of the user who created them – often with high privileges and broad access – and can generate costs in compute usage, token consumption and remediation effort when something goes wrong.

That does not mean organizations are not attempting to develop frameworks and guidance. The CSA reports that 82% of enterprises have unknown AI agents in their environments.¹⁹ The CSA has established working groups around areas such as AI Controls²⁰ and has published a substantial library of guidance including *Autonomous but Not Controlled*, *Enterprise AI Security Starts with AI Agents*, *Identity and Access Gaps in the Age of Autonomous AI* and *Securing Autonomous AI Agents*, among others.²¹ The CSA has also published MAESTRO (Multi-Agent Environment, Security, Threat, Risk and Outcome), an agentic AI threat modeling framework.²²

Analysis, aggregation and reporting standards that provide a single structured language for data from multiple sources are an important tool for effective publication, collection and reuse of information.

Tavro has offered the OWASP AI Vulnerability Scoring System (AIVSS) as a tool to evaluate a structured and quantifiable methodology to identify, assess and mitigate vulnerabilities specific to all types of AI systems, including agentic AI. OWASP -- the Open Worldwide Application Security Project – is a nonprofit foundation working to improve the security of software.

How Do You Design an Agent Metadata Specification?

As with any kind of data gathering, analysis, aggregation and reporting, standards that provide a single structured language for data from multiple sources are an important tool for effective publication, collection and reuse of information. Those of us in Extensible Business Reporting Language (XBRL) circles have long promoted this principle – providing a common structured

¹⁸<https://cloudsecurityalliance.org/artifacts/enterprise-ai-security-starts-with-ai-agents>.

¹⁹<https://cloudsecurityalliance.org/press-releases/2026/04/21/new-cloud-security-alliance-survey-reveals-82-of-enterprises-have-unknown-ai-agents-in-their-environments>.

²⁰<https://cloudsecurityalliance.org/research/working-groups/ai-controls>.

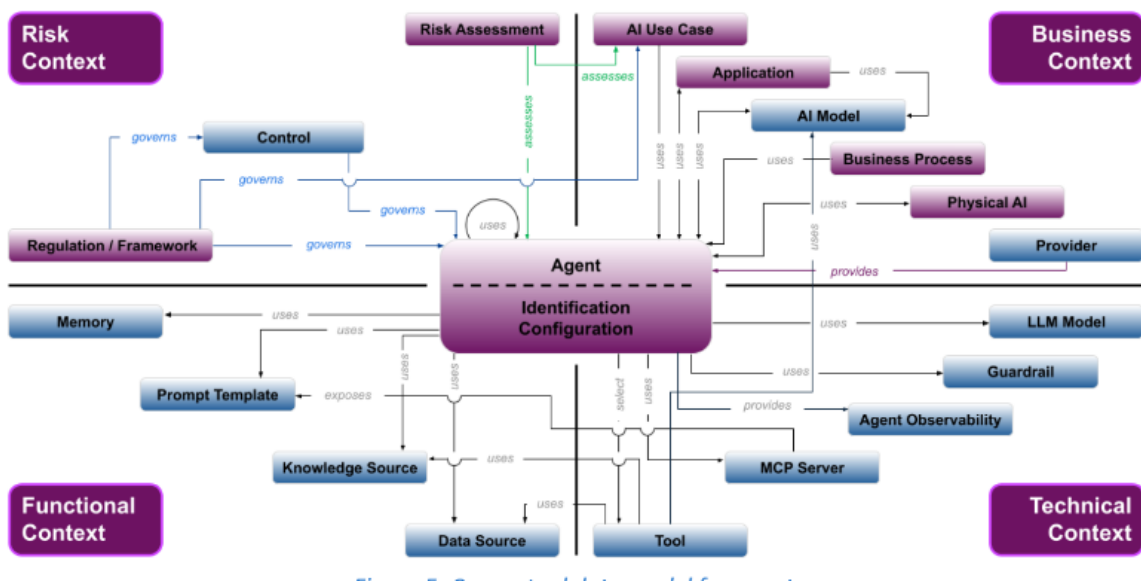
²¹<https://cloudsecurityalliance.org/research/publications?term=artificial-intelligence>.

²²<https://cloudsecurityalliance.org/blog/2025/02/06/agentic-ai-threat-modeling-framework-maestro>.

language that enables consistent data from multiple sources to be compared and acted upon. I have pursued this vision of a "single source of truth," a holistic view of the organization, through XBRL GL (Global Ledger), a standard designed to represent detailed transactional data across systems.

For this reason, the Agent Metadata Specification sets out to provide consistent agent metadata to support agent risk management at scale. Tavro's conceptual model, illustrated in Figure 1, organizes the specification across four key contexts considered necessary to provide a consolidated, holistic view of all deployed agents across an entire organization.

Figure 1: Tavro's Agent Metadata Specification Conceptual Model
 Source: <https://github.com/TavroOrg/Agent-Metadata-Specification>



By implementing this robust metadata specification, enterprises gain several critical advantages: a unified enterprise-wide view and single source of truth for all deployed agents; enhanced accountability and transparency of ownership, with the metadata structure capturing ownership details that ensure clear accountability for agent performance, maintenance, and policy adherence; automated risk management functions across all business processes and the agents they consume; and significantly accelerated audit readiness for governance and compliance purposes.

What the Author Added: Extending Governance to Human Actors

When I first encountered Tavro's work, the specification struck me as broadly a good idea – but incomplete. The numbers quoted by the CSA and others illustrate that there is indeed an AI Agent Visibility Vacuum. The documents published by Tavro speak eloquently about the potential benefit of better control over AI agents and their activities. But one thing stood out as a major missing element: we do not have tools to track *human* activities in the manner we hope

to track AI agents, nor do we have tools designed to track the interactions between humans and the AI agents they direct, supervise, or delegate to.

I have, for a quarter of a century, encouraged the accounting profession to leverage digital signatures for human authorization, authentication, and activity capture – not as a surveillance mechanism, but to create an electronic audit trail of responsibility and activity. In 2026, we still do not have a single, holistic view of important acts by individuals across systems and platforms. The Agent Metadata Specification offers an opportunity to remedy that.

The Governance Symmetry Between Humans and AI Agents

As I documented in my GitHub repository, just about everything true about AI agents as governance subjects is also true about people. AI agents and human actors share a common set of governance requirements:

- Both take actions that affect systems and data.
- Both can hold privileges, scopes, and access rights.
- Both can cause harm, errors, or fraud.
- Both need to be monitored for compliance and auditability.
- Both operate across multiple systems and platforms.
- Both require identity, provenance, and accountability metadata.

By unifying human and agent metadata under a common specification, organizations can bridge the gap between traditional identity and access management systems and the new world of autonomous AI agents - creating a single, coherent audit trail that spans all actors in a workflow.

The Author's Contribution to the Specification

The author's fork of the specification²³ proposes extending the metadata model with an explicit *Actor* abstraction - a generalization that encompasses both AI agents and human actors under a shared governance schema. The key additions include:

- **Human actor identity metadata.** Analogous to agent identity fields (such as agent ID, name, version, and deployment context), the specification gains fields for human actor identification: user identity, role, organizational unit, and credential or authentication method.
- **Provenance and delegation chain metadata.** Fields to record the originator of an instruction or task, the sequence of delegations that followed, and the ultimate actor (human or AI) that executed each step.
- **Interaction event metadata.** A record of human-agent interactions: prompts issued, outputs reviewed, approvals granted, overrides exercised, and escalations triggered.
- **Scope and permission alignment.** Fields to capture the permissions under which an actor operated at the time of action - enabling after-the-fact comparison of intended scope with actual scope of activity.

Why This Matters: The Delegation Chain in Practice

²³ <https://github.com/GenAICPA/Agent-Metadata-Specification>

Consider the period-end close scenario introduced at the opening of this article, viewed now through a governance lens. A controller (human) issues a natural language instruction to Agent A, tasking it with closing the books for the period. Agent A assesses the work required and delegates foreign currency translation to Agent B and accounts receivable reconciliation to Agent C. Agent B executes a set of journal entries in the ERP system; Agent C queries the accounts receivable subledger, flags three disputed items, and logs exceptions.

Without unified actor metadata, the audit trail of this process is fragmented: the controller's instruction exists in a chat or email interface; Agent A's delegation decision exists in an agent orchestration log; Agents B and C's actions exist in separate system logs, if they are logged at all. No single record connects the human instruction to the downstream machine actions, nor establishes that each step was performed within the intended scope of authority.

With unified actor metadata, the reconstruction is straightforward. Each step in the chain - Human → Agent A → Agent B → System Action; Human → Agent A → Agent C → Exception Log – is recorded with a common schema capturing actor identity, provenance, delegated scope, timestamp, and outcome. An auditor can traverse the chain in either direction. A risk system can flag deviations from expected scope in real time. Regulators can request a full record with confidence that it is complete.

Governance frameworks that hold AI agents to a higher standard of traceability than the humans who direct them are structurally incomplete.

Governance frameworks that hold AI agents to a higher standard of traceability than the humans who direct them are structurally incomplete. Accountability does not improve by making only one half of the actor ecosystem visible. As the author's specification states, the goal is "symmetry in accountability" – event-level traceability for all actors, regardless of whether they are human or machine.

Where Do We Go from Here?

I was given the opportunity for an initial meeting with the Tavro team, and their work has continued to advance. As of this writing, my fork is some 50 commits behind the main repository – a sign of the momentum the project has gained since we last synchronized. That gap is an invitation, not a barrier. The specification is an open-source project, and the governance problems it addresses are real, growing, and shared across industries.

For accounting and financial professionals – and the governance, risk and compliance practitioners who serve them – the emergence of AI agents represents both a new category of risk and a new tool for assurance. The following four areas represent a practical framework for engagement:

Assess the Efforts

Accounting, audit, and risk professionals should review both the Tavro Agent Metadata Specification and the author's extension covering human and human-AI actor interactions. These efforts reflect the leading edge of thinking about how to govern an environment in which AI agents operate alongside – and increasingly, on behalf of – human actors. The GitHub

repositories are open, the documentation is accessible, and the community welcomes practitioners with domain knowledge in governance, risk, and compliance.

Identify Reporting Frameworks and Internal Risk Needs

Organizations should assess how an agent metadata standard maps to their existing internal risk frameworks, control environments, and reporting obligations. Which business processes now involve AI agents? Where do agents interact with human decision-makers? What access and privilege levels do agents currently operate under, and are those levels documented, reviewed, and appropriate? The answers to these questions define an organization's governance gap – and the metadata specification provides a structural language for closing it.

Established reporting frameworks – XBRL-based disclosures, SOX control narratives, internal audit programs – will need to evolve to reflect agentic workflows. The period-end close, the vendor payment process, the financial statement review: wherever a human-in-the-loop assumption has been embedded in a control, that assumption now requires examination.

Support Emerging Governance Efforts

The Cloud Security Alliance, the Open Worldwide Application Security Project, and Tavro are among the organizations actively building the governance infrastructure that enterprises and regulators will rely on. The accounting and audit profession has both the credibility and the domain knowledge to contribute meaningfully – particularly in areas such as assurance standards for AI-driven processes, internal control frameworks adapted to agentic environments and the standards for human-agent accountability that this specification begins to address.

The EU AI Act has created regulatory momentum in Europe, and analogous frameworks are developing globally. The accounting and audit profession's involvement in shaping these frameworks – as it has shaped financial reporting standards for decades – is both an opportunity and a responsibility. The time to engage is before standards are settled, not after.

Monitor and Contribute to the Actor and Agent Metadata Specification Work

The Agent Metadata Specification is a living document. Both the Tavro repository (<https://github.com/TavroOrg/Agent-Metadata-Specification>) and the author's extension (<https://github.com/GenAICPA/Agent-Metadata-Specification>) are open to review and contribution. Practitioners who manage agents in production, who conduct audits of AI-driven processes, or who design governance and risk frameworks are precisely the contributors the specification needs. The technical community can define the schema; the practitioner community must validate that the schema captures what governance actually requires.

The age of agents arrived quickly, broadly, and without waiting for the governance infrastructure to catch up. Accounting and financial professionals have a long tradition of building the frameworks that make complex, distributed activity legible, auditable and accountable. The challenge of governing AI agents – and the humans who direct them – is the latest expression of that tradition. The tools are being built. The profession should be at the table.





ThinkTWENTY20

The Magazine for Financial Professionals

Meet a Very Special Group of People.....

Our readers are financial professionals – mostly accountants – who have a thoughtful interest in the issues shaping their profession. They read our magazine because it provides more substance than most other sources. They are members of all the major firms, regulators, business people and academics.

They are also interested in learning how *your* company can help them address the issues confronting them, and how your services and products can transform their work and workplace to meet the future head on.

Many of the profession's top thought leaders are in our audience and CPA Canada, Canada's largest professional accounting organization, is a Founding Partner of **ThinkTWENTY20**.

If you have products or services aimed at the accounting profession, you have several ways to reach out to this elite group. Your advertisement can appear in a premium spot in **ThinkTWENTY20** or on our webpage, or both. You can also tell your story in a sponsored article in either the magazine or on our webpage, or both – explaining how your products or services can enhance productivity, safeguard data, improve security, streamline or automate processes, etc. Our readers tend to pay attention and listen to good advice. They will be interested in what you have to say.

Rates and Specs:

Advertise on our website:

Half page for six months: \$ 450.00
Regular, for six months: \$ 375.00

Advertise in ThinkTWENTY20:

Full page: \$ 450.00
Half page: \$ 300.00

- 30% discount if your ad runs in both the magazine and on the webpage.
- All prices are in Canadian dollars.

Sponsored content: \$500.00 per article. Articles must be approved by our editors.

ThinkTWENTY20 means ThinkCHANGE.



ThinkTWENTY20

The Magazine for Financial Professionals

CHECK OUT OUR E-BOOKS

Recent additions to our published e-books include:

Integrated Thinking and Corporate Social Responsibility. This book sets out a model for companies to follow when they are interested in pursuing strategic initiatives to address environmental concerns. It acknowledges the merits of sustainability reporting but goes beyond that to the idea that such reporting is worthwhile and may point the way for certain strategic initiatives, but that reporting on sustainability-related matters is not enough. The deteriorating state of our planet and the advancement of climate change are so severe that more direct action is needed. It's available on Amazon and in our store.

- [Interviews by ThinkTWENTY20](#)
- [Mental Health in The Accounting Profession – A Major Study](#)
- [Integrated Thinking and Corporate Social Responsibility by Gerald Trites](#)
- [Selected Writings of Eric Cohen on AI.](#)

Available in our store at this link: www.thinkttwenty20.com/shop

Also available on Amazon.