HUNTON

Artificial Intelligence

2025 Year in Review

Capital Markets and Securities

SEC and AI: What Companies Should Learn from 2025 to Prepare for the 2026 Reporting Season

By Mayme Donohue

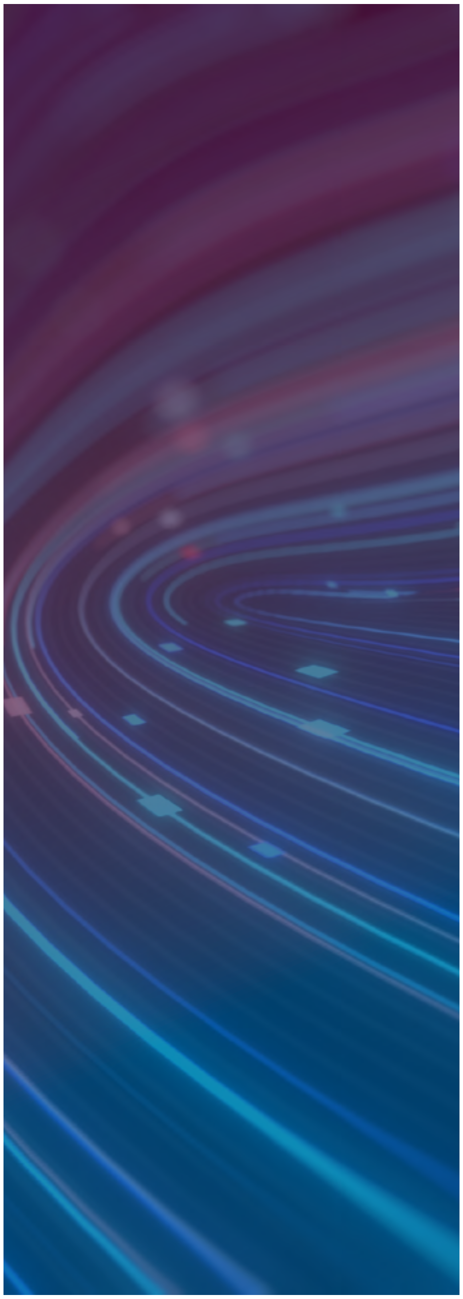

2025 was a transition year in the SEC's posture toward artificial intelligence (AI). The Commission continued to signal that "AI washing" and other AI-linked misstatements remain classic enforcement targets while it leaned into AI internally to modernize its own operations. We saw the Staff's disclosure review program probe AI-related narratives through targeted comments and the SEC clearly messaged that AI would be an examination focus for 2026.

The SEC's 2025 Message on AI Focused on Existing Securities Principles

Across speeches and public statements in 2025, a consistent theme emerged. AI does not require a new set of investor-protection principles to trigger SEC scrutiny. Instead, the SEC repeatedly framed AI issues as variations on familiar securities law concepts:

- **Accuracy and completeness** of statements, particularly around "AI-enabled" products, revenue drivers, competitive differentiation, and R&D claims;
- **Reasonable basis and substantiation** for AI-related assertions, including performance, automation, "proprietary models," and the role of humans; or
- **Material risk disclosure** where AI meaningfully affects operations, cybersecurity, data use, IP, regulatory exposure, or human-capital impacts.

These themes were reinforced late in 2025 when SEC leadership highlighted an Investor Advisory Committee (IAC) workstream on AI disclosure. At the IAC meeting in December, Chairman Atkins reinforced that the "principles-based rules were intentionally designed to allow companies to inform investors of material impacts of any new development, including how AI affects their financial results, how AI can be a material risk factor to an investment, and how AI is a material aspect of their business model." Additionally, Commissioner Uyeda emphasized that the SEC has proposed, at a minimum, that issuers define what they mean by "AI," describe board oversight (if any), and separate discussion of AI's impacts on internal operations vs. customer-facing products and services.

Practical Point for Drafting 10-K/20-F/Registration Statement Disclosures

In 2025, the SEC's posture strongly suggested that "AI" is not a safe buzzword, rather it is a potentially material disclosure that should be treated accordingly within the existing framework of materiality-based disclosure principles.

The SEC Embraced AI Internally with AI Task Force and Chief AI Officer

One of the clearest 2025 developments was institutional when the SEC announced the creation of an AI Task Force to centralize and accelerate responsible AI integration across the SEC, with an emphasis on governance and lifecycle management. The SEC also publicly identified its Chief AI Officer as leading the task force, underscoring that the initiative is meant to be durable and cross-divisional rather than ad hoc experimentation. Separately, the SEC maintained an Artificial Intelligence at the SEC landing page that highlights internal planning, including the SEC's 2025 AI Compliance Plan aligned with OMB AI guidance.

Why This Matters for Issuers

The Commission's internal build is not just operational, it is also a signal that AI governance and controls are becoming table stakes. As the SEC adopts AI-enabled tools, its expectations for how registrants manage similar risks (like data provenance, human oversight, testing/validation, vendor management and documentation) are likely to become more concrete in exams, comment letters, and enforcement.

"AI Washing" Remains an Enforcement Focus

The SEC's messaging in 2025 continued to highlight AI-related misstatements as a priority area. Often labeled "AI washing" (i.e., overstating or mischaracterizing AI capabilities), the SEC kicked off 2025 by settling charges against Presto Automation Inc., a restaurant-technology company that was listed on the Nasdaq until September 2024, for making materially false and misleading statements about critical aspects of its flagship AI product, Presto Voice.

Additionally, the SEC staff issued AI-related comment letters in 2025, including requests for more detail on development, validation, third-party dependencies, and the real operational role of AI/ML. Examples from publicly available correspondence show staff asking companies to expand and operationalize AI-related discussions, for example describing governance policies around AI use, or revising business/risk factor disclosure to more fully address the state of AI adoption and regulatory landscape.

Practical tips to Avoid Inadvertent AI Washing

The risk is not only in investor decks or marketing pages. It can show up in:

- **Business descriptions** that portray AI as core to differentiation without describing the actual state of deployment;
- **Risk factors** that acknowledge generic AI risks but do not align with how the company truly uses data/models/vendors;
- **MD&A** narratives that attribute efficiencies or margin expansion to AI without a clear basis; or
- **Forward-looking claims** about "AI roadmaps" that are inconsistent with budget, staffing, vendor contracts, or product readiness.

AI Is on the List of 2026 Examination Priorities

The SEC's Division of Examinations identified AI as a focus area in its Fiscal Year 2026 Examination Priorities, emphasizing that it will be analyzing registrant's AI-related disclosures focusing on "recent advancements in AI and will review for accuracy registrant representations regarding their AI capabilities." The SEC is not hiding the ball, and combined with the public statements from the Chair and other commissioners along with the 2025 comment letter trends, companies should not take their AI-related disclosures lightly.

2026 Practice Pointers for AI-Related Public Disclosures

- **Inventory and Map AI Use Cases** Identify where AI/ML is used across the business (product, operations, finance, HR, cybersecurity, compliance, legal, customer service) and separate pilot, internal-only, third-party enabled, and customer-facing uses.
- **Pressure-Test External Statements** Validate claims in earnings scripts, roadshow decks, investor presentations, web copy, and product collateral to confirm that "AI-enabled" statements reflect real functionality and not marketing shorthand.
- **Align Risk Factors to the Company's Actual AI Profile** Consider topics like data provenance and usage rights; IP risks (training data, outputs, open-source/model licensing); cyber and fraud risks (including deepfakes and social engineering); and regulatory exposure (sector-specific rules, cross-border data regimes).
- **Evaluate Governance and Disclosure Controls** Document oversight (board/committee, management steering group, escalation paths), implement vendor and model risk management (testing/validation, monitoring, change management) and treat AI-related disclosure as a disclosure-controls topic, not just "innovation messaging."

Privacy and Cybersecurity

CCPA Automated Decision-Making Technology Requirements

By Michael La Marca and Raul Orozco

In 2025, California added to the growing patchwork of state rules governing the use of automated decision-making technology (ADMT). Several states now have comprehensive privacy laws that include restrictions on profiling (i.e., any form of automated processing of personal data to evaluate, analyze, or predict certain personal aspects about an individual) in furtherance of decisions that produce legal or similarly significant effects concerning a consumer. In addition, there is an emerging body of sectoral rules regulating the use of ADMT in certain industries, including in areas such as employment, insurance, housing, and healthcare. Relatedly, the Colorado AI Act, which was enacted in 2024 and is due to become effective in July 2026, will regulate the use of "high-risk AI systems" that are used to make certain types of "consequential decisions" concerning Colorado residents. While these rules differ in scope and application, they often share a number of commonalities, including transparency obligations, governance requirements (e.g., policies and impact assessments), and individual rights (e.g., the right to access information regarding the ADMT and the decision and the right to opt out of and/or appeal the automated decision).

In September 2025, California's Privacy Protection Agency (CPPA) finalized regulations pursuant to the California Consumer Privacy Act (CCPA) that will govern the use of ADMT.

- ADMT is defined as "any technology that processes personal information and uses computation to replace … or substantially replace human decision-making."
- ADMT "substantially replaces human decision-making" when a business uses its output without "human involvement," i.e., a human reviewer that (1) knows how to interpret and use the technology's output to make the decision, (2) reviews and analyzes the output of the technology, and any other information that is relevant to make or change the decision, and (3) has the authority to make or change the decision based on their review and analysis.
- Although initial drafts of the ADMT regulations had a broader scope, the finalized regulations apply to the use of ADMT to make "significant decisions" concerning California residents, meaning decisions that result in the provision or denial of financial or lending services, housing, education enrollment or opportunities, employment or independent contracting opportunities or compensation, or healthcare services.

Article 11 of the updated CCPA Regulations establishes a number of new obligations for businesses that use ADMT to make significant decisions concerning California residents. These requirements, which are discussed below, will become effective on January 1, 2027.

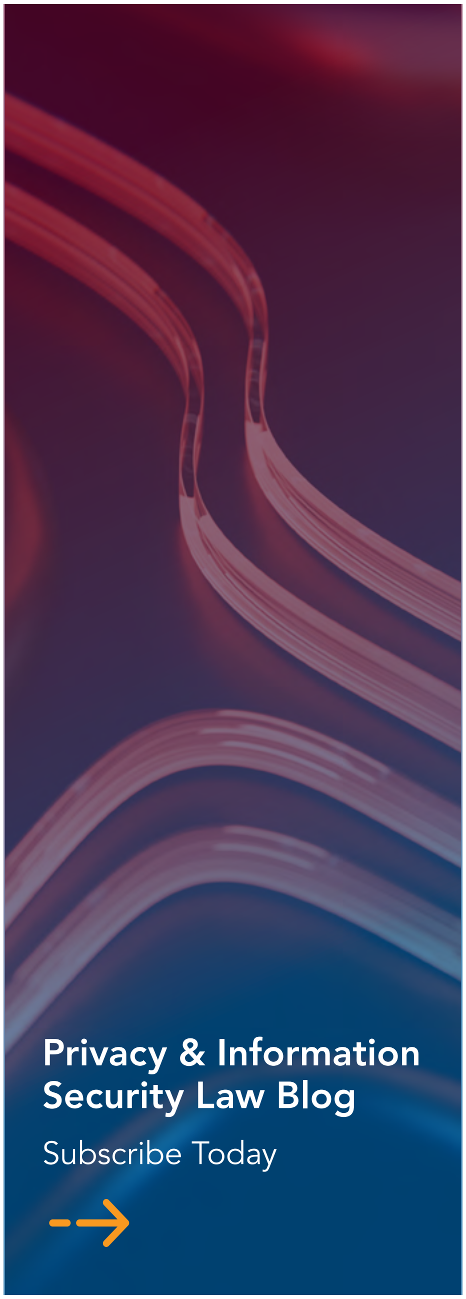

Pre-Use Notice Requirements

Article 11 of the updated CCPA Regulations requires businesses that use ADMT to make significant decisions to provide California residents with a "pre-use notice" describing the intended use of ADMT and the individual's related rights. The notice, which must be presented prominently at or before the point of collecting personal information that will be processed using ADMT, must:

- Explain the specific purpose for using the ADMT (a generic description, such as "to make a significant decision" is insufficient);
- Describe the consumer's right to opt out (or, where an exception applies, the right to appeal to a human reviewer);
- Explain the consumer's right to access certain ADMT-related information;
- Include a plain-language description of how the ADMT operates, including the personal information used, the types of output generated by the ADMT, how those outputs inform decisions, and how decisions would be made by the business if a consumer opts out of the ADMT; and
- State that the business may not retaliate against consumers for exercising CCPA rights.

Given these requirements, businesses that use ADMT to make significant decisions will need to ensure that the technology is sufficiently explainable (i.e., it is possible to determine both how the technology produces results at a general level and why a particular decision was reached), which can be challenging given the "black box" nature of many proprietary AI models. In addition, businesses will need to ensure they have personnel with adequate expertise and training (often referred to as "AI literacy") to understand the underlying logic and functioning of the ADMT and translate that understanding into a plain-language description for a layperson audience.

Importantly, businesses are permitted to exclude from their pre-use notice information that constitutes trade secrets or that would compromise the business's ability to address security incidents, fraud, or other illegal activity. Businesses will thus need to strike a balance of addressing the CCPA's emphasis on transparency while refraining from divulging certain proprietary or security-sensitive information, particularly when explaining how the ADMT operates.

Consumer Rights to Opt Out of and Access ADMT

Article 11 grants California residents the right to opt out of a business's use of ADMT to make significant decisions. Businesses must provide at least two easy-to-use opt-out methods, one of which must reflect the business's primary mode of interaction with California residents, along with a means for confirming that the business processed the opt-out request. To the extent a business interacts with California residents online, it must provide an opt-out link in its pre-use notice (e.g., Opt Out of Automated Decision-making Technology) that directs individuals to an interactive form where they can opt out.

Businesses are not required to provide California residents the right to opt out if:

- The business provides a method for appealing the decision to a designated human reviewer who (1) has the authority to overturn the decision, (2) knows how to interpret and use the ADMT's output, and (3) in response to a request to appeal, reviews the output of the ADMT and any other information relevant to the decision (including information provided by the individual in support of their appeal).
- The ADMT is used solely for certain employment or educational decisions (e.g., allocation/assignment of work or compensation decisions) and does not otherwise unlawfully discriminate based on protected characteristics.

If a consumer opts out after the ADMT processing has already begun, the business must cease the processing within 15 business days and instruct downstream recipients to do the same.

In addition to honoring opt-out requests, businesses must respond to verified requests to access ADMT-related information, including plain-language explanations of:

- The specific purpose for which the business used ADMT with respect to the individual who submitted the access request;
- The logic of the system, sufficient to enable the individual to understand how the ADMT processed their personal information to generate an output (e.g., an explanation of the parameters the ADMT used to generate the output);
- The outcome and nature of the decision-making process for the individual (including whether the ADMT's output served as the sole factor in the decision and the extent to which a human played a role in the decision); and
- The individual's rights under the CCPA.

Although much of the information that must be provided in response to an access request tracks the requirements of the pre-use notice, a business's response must be specific to the individual making the request, which again highlights the importance of ensuring that the ADMT has sufficient explainability and that the business has the requisite internal expertise to understand and convey how the ADMT reaches a decision. As with the pre-use notice, businesses are permitted to withhold trade secrets and security-sensitive information from an access request.

Risk Assessments

In addition to the new ADMT requirements under Article 11 that are described above, businesses should be aware that requirements relevant to ADMT appear throughout the updated CCPA Regulations. For example, business that use ADMT to make significant decisions about California residents must update their privacy notices to describe the right to opt out of such decisions and access ADMT-related information. In addition, Article 10 of the updated CCPA Regulations imposes a new requirement to conduct detailed risk assessments for certain processing activities that present a "significant risk" to the privacy of California residents, and submit certain information regarding those risk assessments (including an attestation of compliance) to the CPPA. Processing activities that present a significant risk to California residents include:

- Using ADMT to make significant decisions about CA residents; and
- Using (or intending to use) a California resident's personal information to train an ADMT for a significant decision concerning a consumer.

Businesses that use ADMT to engage in significant decisions (or train ADMT to be used for such purposes) should thus review the entirety of the updated CCPA Regulations to ensure they identify and address all relevant compliance obligations.

Conclusion

California's updated CCPA Regulations establish new standards for ADMT, emphasizing transparency, governance, and consumer rights. As ADMT becomes increasingly embedded in critical decisions, and regulatory scrutiny over such technology grows, businesses should ensure they are monitoring the evolving regulatory landscape (including the growing convergence of privacy and anti-discrimination requirements under emerging AI laws) and preparing for heightened compliance obligations. Companies that invest in robust AI governance and risk management programs will stay ahead of the curve in their ability to navigate and adapt to a shifting regulatory environment while minimizing legal and reputational risks.

Intellectual Property

The USPTO’s Revised (2025) AI-Assisted Invention Guidance

By Steven Wood

Close to two years ago—in February 2024—the US Patent and Trademark Office (USPTO) released detailed guidance addressing the issue of whether AI could be named as an inventor on a patent application. 89 Fed. Reg. 10043 (Feb. 13, 2024) (2024 guidance). The short answer was that AI cannot be named as an inventor or joint inventor.

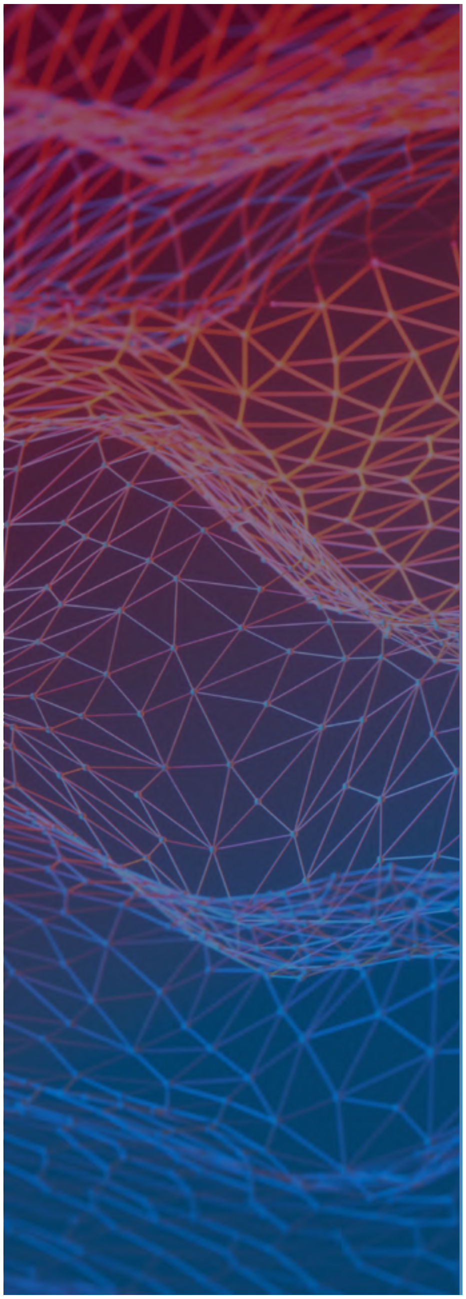

Now, the USPTO has rescinded the 2024 guidance and provided inventors and patent applicants with new guidance, updating the framework regarding AI-assisted inventions and how such will be examined at the USPTO. 90 Fed. Reg. 54636 (Nov. 28, 2025) (2025 guidance).

The 2024 Guidance Focused on a “Significant Contribution” by a Human

The 2024 guidance did not alter the human inventorship requirement for patents: AI cannot be named as an inventor and inventions created *entirely* by AI are not patentable. However, it introduced the concept of whether a human “significantly contributed to the invention,” meaning that a person had to do more than merely rely upon an AI system to come up with an invention in order to obtain a patent on that invention (i.e., be named an inventor). The 2024 guidance relied on application of the *Pannu* factors to determine inventive contribution. *Pannu v. Iolab Corp.*, 155 F.3d 1344 (Fed. Cir. 1998). Recognizing that there is no bright-line test for determining the significance of a contribution, the USPTO included two examples along with supporting analysis to provide some clarity on application of the test for practitioners and inventors.

The 2024 guidance also contained a section on the duty of disclosure detailing how proper inventorship is material to patentability. Improper inventorship means a fraud has potentially been committed on the USPTO, making the patent unenforceable. To that end, the 2024 guidance focused on ensuring that the required duties are met regarding disclosure to the USPTO.

Now, the Human Inventorship Requirement Is the Focus

The new 2025 guidance takes a more straightforward approach.

While the core principle remains that AI cannot be named as an inventor, the perhaps the most critical takeaway from the 2025 guidance is its emphasis that the legal standard for determining inventorship applies to “all inventions, regardless of whether AI systems were used in the inventive process.” This is a change. As noted above, the 2024 guidance focused on the use of AI systems in the inventive process. Instead, the 2025 guidance cites *Thaler v. Vidal*, 43 F. 4th 1207 (Fed. Cir. 2022), which holds that only natural person(s) can be named inventors and centers the inventorship inquiry around “conception” as the “touchstone of inventorship,” a fact intensive inquiry.

Regarding the *Pannu* factors, the 2025 guidance notes that the USPTO presumes the inventor(s) presented on the patent application are the correct inventor(s). AI is described as a mere tool that can be used by an inventor. While the conception inquiry standards apply, the *Pannu* factors (to determine inventive contribution) only apply if joint inventorship (of humans) is at issue. With a sole inventor, the factors do not apply because AI cannot be a joint inventor, so no analysis is required.

Lastly, the 2025 guidance notes that the inventorship standard applies also to design and plant applications in the same manner as utility patent applications.

And, in contrast to the previous guidance, there is no discussion of the duty of disclosure in the new guidance.

Priority Claims to Foreign Applications Cannot Include AI as a Named Inventor

Given the USPTO’s position that AI cannot be named as an inventor, this seems like common sense, but the final section of the 2025 guidance specifically calls out that:

> Applications and patents claiming the benefit of, or priority to, a prior application filed in the United States or a foreign country under 35 U.S.C. 119, 120, 121, 365, or 386 must name the same inventor or have at least one joint inventor in common with the prior-filed application.

The guidance goes on to caution that a priority claim to a foreign application that names an AI as the sole inventor will not be accepted, noting this policy also applies to US patent applications and patents claiming priority to foreign applications that allow the naming of non-natural persons as joint inventors.

But among foreign patent systems, currently only South Africa allows the naming of AI as an inventor. Thus, the inclusion of this section of the 2025 guidance is interesting and may be forward leaning to account for additional patent systems (in other countries, in the future) allowing the naming of AI as an inventor or joint inventor on patent applications.

The Human Element Tracks with Copyright Authorship Guidance

The USPTO’s revised 2025 guidance is consistent with the position of the US Copyright Office, as well as courts, that AI alone cannot be an author for purposes of copyright law, but the mere presence of AI in the creation of a work does not doom a copyright application. Similarly, any material created by AI must be disclosed (if more than *de minimis*) and disclaimed by the human author (i.e., AI generated material is not protectable by copyright). This affirms that the creative expression for which copyright is sought stems from the human mind.

Takeaways

The key takeaway for patent applicants is that humans remain central to inventorship of patentable ideas. If AI is used in the invention process, it is only a tool or instrument (in the eyes of the USPTO). The inventorship inquiry centers on the actual persons involved, not the tool—which is the way it previously was (that is, prior to the 2024 guidance that possibly blurred the issue).

As a best practice, documentation is important to inventorship, regardless of whether AI is involved in the inventive process leading to a patent application. While the USPTO may not challenge inventorship, subsequent litigation could expose flaws. Accordingly, documentation of each step in the patenting process (e.g., with lab or inventor notebooks) is fundamentally important to ensure that the record is clear on the role of each person involved with the inventive process (as well as any AI that was used).

And finally, even though not explicitly discussed in the 2025 guidance, the duty of disclosure from 37 CFR 1.56 has not changed, along with other duties of a practitioner and those associated with the prosecution of patent applications before the USPTO.

Labor and Employment

AI in HR: Year in Review

A Review of 2025 with an Eye Toward 2026 and Beyond

By Robert Dumbacher and Daniel Butler

Though Congress did not pass any laws in 2025 that specifically targeted employers' use of AI in the workplace, legislative development continued at the state level. Indeed, several states passed AI legislation or issued regulations that have taken or will take effect in 2026 and 2027. As a result, employers with operations in multiple states should carefully review AI legislation at the state level to ensure compliance.

Looking ahead, employers also will want to monitor a possible conflict between the patchwork of state laws and federal executive action. On December 11, 2025, President Trump issued an Executive Order entitled "Ensuring a National Policy Framework for Artificial Intelligence," that promises to, within 30 days, establish an AI Litigation Task Force to evaluate state laws on AI and potentially challenge any laws "inconsistent" with federal policy to "sustain and enhance the United States' global AI dominance through a minimally burdensome national policy framework for AI."

Significant State Level Developments in 2025

Arguably, the most significant state law developments in the AI-HR space for 2025 occurred in California and Colorado.

California

On June 30, 2025, California's Civil Rights Council issued regulations to clarify that the state's existing anti-discrimination protections under the Fair Employment and Housing Act (FEHA) apply when employers (or their vendors) use AI, algorithms, or other "automated-decision systems" in employment. These new clarifying regulations are already in place, with an effective date of October 1, 2025.

Separate from California's anti-discrimination law, the California Privacy Protection Agency (CPPA) finalized privacy rules for automated decision-making technology (ADMT). These rules include transparency obligations, opt-out rights to employees for certain high-impact uses, and risk-assessment duties—covering "significant decisions" such as hiring, promotions, compensation, and termination. These new rules will go into effect on January 1, 2027. Employers using AI tools in California would be wise to begin implementing changes to come into compliance with the new rules in 2026. These steps should include, at minimum, employee/candidate notices, risk assessments, establishing access/explanation rights, and documenting assessments with human oversight.

Colorado

Colorado made headlines through much of 2024 and 2025 with its Senate Bill 24-205, the Colorado Artificial Intelligence Act. Though the law's effective date has been delayed until the end of June 2026, the penalties for noncompliance can be steep—up to $20,000 per violation. Though there is no private right of action under the law, the Colorado Attorney General's office has jurisdiction to enforce compliance. Generally, the law will require developers and deployers (e.g., employers) of "high-risk" AI systems to use reasonable care to prevent algorithmic discrimination and follow a governance playbook: risk-management programs, impact assessments, annual reviews, consumer (candidate/employee) notices, correction rights, and human appeal of adverse decisions. The statute and official summaries make clear that employment decisions are squarely in the scope of "high-risk" users of AI systems.

To comply, employers should lean on the law's safe harbor provisions, which mandate a rebuttable presumption of compliance if a deployer of a high-risk AI system adopts a number of practices, including a risk management policy governing AI systems, annual impact assessments, informing individuals when they are interacting with an AI system, and posting a website notice about the use of AI systems.

December 2025 Federal Executive Order and Possible Future Preemption

At the federal level, Congress has proposed several laws around AI and the workplace, such as the "No Robot Bosses Act," but to date, no bills have been enacted. Of note, existing federal civil-rights law applies to AI (for example, disparate-impact and disparate-treatment theories under Title VII of the Civil Rights Act of 1964). And a 2024 court decision in California held that disparate impact claims under Title VII could, in certain circumstances, be made against AI vendors, as agents of employers.

Looking ahead to 2026 and beyond, a potential conflict may be brewing between federal enforcement authority and state laws. The President's December 11, 2025, Executive Order seeks to establish a uniform framework for AI governance across the nation. To accomplish that objective, the EO directs the Attorney General to study state AI-laws and create a litigation task force that will challenge AI laws inconsistent with federal policy. The stated federal policy in the EO is somewhat vague but emphasizes "minimally burdensome" regulation that ensures United States' "global AI dominance." The EO also tasks the special advisor for AI and crypto and the assistant to the president for science and technology to prepare a legislative recommendation that would establish a uniform federal policy framework for AI that "preempts state AI laws that conflict with the policy set forth in this order."

Employer Action Items Today

For now, employers should comply with applicable state laws already implemented or that are scheduled for implementation within the coming year. To do this, employers should generally, at minimum, identify every AI tool that touches upon employment decisions, and determine whether those tools might be regulated by a state or local specific legal regime. If necessary, employers should perform bias audits, at least annually, to ensure their AI tools are not violative of disparate impact principles. A sound AI policy also includes robust internal policies surrounding notice of AI uses, employee access to AI data, and mandated human oversight of AI recommendations. Partnering with counsel can help promote these decisions being protected by attorney-client privilege.

Looking ahead, employers should also keep abreast of continued state and local legislative and regulatory actions, as well as potential federal enforcement actions and congressional developments.

Insurance

Artificial Intelligence and Insurance

By Michael S. Levine, Alex D. Pappas, Casey Coffey, and Madalyn Moore

As outlined earlier this year, insurance policies, including directors and officers (D&O) policies and commercial general liability policies, should provide coverage for artificial intelligence-related losses. The reason for this is simple: many of the risks posed by AI are the same as those that have long been covered by standard-form policies. But as the year progressed, the insurance landscape for AI-related risks changed as insurers introduced AI-specific exclusions to their traditional lines of coverage. With that, the year also saw the continued launch of additional new products that specifically target AI risks. These changes signal a new era for AI-risk management.

The Rise of AI Exclusions in 2025

2025 witnessed a proliferation of policy wording changes specifically attempting to limit or exclude coverage for AI-related losses. Below are a few examples of the newer provisions.

Insurers	Exclusions
Hamilton Select Insurance Inc.	Hamilton introduced an exclusion in certain professional liability policies that removes coverage for claims arising out of the actual or alleged use of generative artificial intelligence by the insured.
Insurance Services Office (ISO)	ISO introduced artificial intelligence exclusions for use in commercial general liability (CGL) policies that purport to exclude coverage for claims involving bodily injury, property damage, personal injury and advertising injury that arise out of the use of AI.
Lloyd's of London	The London Market introduced an exclusion that attempts to bar coverage for bodily injury, property damage, or economic loss caused by or resulting from the actions or decisions of artificial intelligence systems.
Philadelphia Indemnity Insurance Company	Philadelphia introduced an exclusion that applies to offenses committed by the insured that are created using generative artificial intelligence in performance.
The Cincinnati Insurance Company	Cincinnati introduced an AI exclusion in D&O liability policies that excludes coverage for the development, deployment, or use of artificial intelligence.
Berkley Insurance Company	Berkley introduced a so-called absolute AI exclusion, which purports to broadly bar coverage for any actual or alleged use or development of AI by anyone.

Policyholders should be vigilant about reviewing their renewal or newly incepting policies for AI-related provisions and consider the coverage gaps these provisions may create. Equally important, policyholders must take inventory in how their company is using AI and assess the risks that AI poses based on its particular uses. No two companies are the same when it comes to how and where they are using AI, and with the rapid deployment and evolution of the technology, even the most robust AI assessment today will likely be obsolete before the end of the current policy period, necessitating a continual evaluation of AI use and risk assessment.

New Affirmative AI Insurance Products

To address the gaps that have been (or will be) created by AI-specific exclusions, insurers have begun to launch new AI-specific insurance products. We highlight several below.

Insurers	Exclusions
AXA XL	AXA XL added AI-specific coverage to its existing cyber insurance products to address the risks of data poisoning, copyright infringement, and liability resulting from the European Union's AI Act.
Chaucer Group and Armilla	Chaucer, a subsidiary of China Re., and Armilla, and AI start up insurer, have joined forces to launch a product that addresses third-party liability for AI system failures.
Coalition	Coalition added AI-specific endorsements with enhanced AI-specific coverage to existing cyber insurance products.
Google	Google announced a partnership with several insurance companies to offer AI-specific insurance options for Google Cloud customers.
Munich Re	Munich Re launched aiSure, which offers performance guarantee coverage for AI technologies, protecting against failures to meet expected results.
Relm	Relm offers insurance products like PONTAAI, which offers coverage for damages, claims expenses, and civil fines (where applicable) arising out of negligent acts, errors, or omissions in AI services.
Testudo	Testudo introduced insurance offerings focused on the unique exposures of AI-driven technology companies.
Vouch	Vouch launched AI insurance.

These new AI products may provide a solution for policyholders who are looking to protect against AI-risks. As with any insurance, however, policyholders should carefully review their particular policy language and work with coverage counsel to ensure the product adequately protects against the risks that they are looking to insure.

What Policyholders Can Expect in 2026

As the risks associated with the use of AI continue to crystallize, policyholders can expect insurers to increase their development and implementation of AI-specific exclusions and limitations. Policyholders also can expect to find a greater variety of AI-focused insurance products to fill the void that exclusions and limitations might create in legacy lines of coverage. Policyholders should be proactive to safeguard against AI-related liabilities. This includes conducting comprehensive and frequent assessments of the ways their company is using AI, staying informed about the evolving regulatory landscape, and regularly reviewing and updating their insurance policies to ensure alignment with specific AI exposures. Collaborating closely with experienced brokers and coverage counsel will be essential to understanding new policy terms and ensuring that appropriate insurance and other risk transfer protection is in place to protect their organization's unique AI risk profile. By anticipating potential challenges and acting strategically, policyholders can strengthen their risk management frameworks and confidently navigate the complexities of AI adoption in the coming year.

Banking and Finance

2026 Top Five Agentic AI Issues in Banking and Financial Services

By Erin Fonté

As we enter 2026, numerous developments in artificial intelligence technology (such as agentic AI) will drive key strategy issues and key decision points for banks and credit unions looking to future-proof products and services for their customer base, as well as exploring new opportunities. Here are five key financial services technology developments and areas to watch in 2026.

How Financial Institutions Will Use AI

Banks spent time in 2025 identifying potential use cases for artificial intelligence and developing policies to both guide its use within organizations and to meet legal and regulatory compliance obligations. But 2026 will see an increase in financial institutions having to address AI questions and issues for the organization's internal use. In the absence of federal regulation, state legislators have passed laws focused on transparency, discrimination, and AI's potential for consumer harm. On December 15, 2025, President Trump issued an executive order to curb state-level actions and to work with Congress to pass a "minimally burdensome national standard"—one that would undoubtedly affect the banking sector. Meanwhile, states (such as California, Colorado, Florida, and Texas, to name a few) are still pursuing state-level AI laws and regulations. So, this is a fast-moving issue.

On the proactive side, financial institutions will need to determine what role AI will play in the organization's overall strategic planning and roadmap. How do financial institutions prepare for this journey while meeting legal, regulatory, and compliance obligations?

- **Four Fundamentals Driving Internal Strategy:** (1) reimagining the customer experience (increased personalization, "frictionless" journeys, etc.); (2) using AI to augment human decision making; (3) modernizing bank/credit union core technology; and (4) setting up a "platform" operating model for the bank's products and services.
- **Banks Using AI Effectively Tend to Focus on Four Global Areas:** (1) setting a bold, enterprise-wide vision for the value AI can create; (2) transforming entire domains, processes, and journeys versus focusing on narrow AI use cases only; (3) building a full AI stack, increasingly powered by multiagent systems; and (4) sustaining and scaling value by setting up critical enablers of AI transformation.
- **Orchestrating Multiagent Systems:** Financial institutions are focused on using multiagent systems to create internal value by automating complex decisions and workflows using AI. Over time, financial institutions could have hundreds of AI agents at their disposal, each trained to complete a particular task and be ready to be called on by other agents or humans. For example, in preparing credit memos, such multiagent systems could yield productivity gains of 20 percent to 60 percent, and faster decision making by 30 percent. This could form the basis of more engaging experiences for customers and financial institution employees.

How Agentic AI in Digital Commerce Will Affect Financial Institutions

On the reactive side, financial institutions will need to understand how to address consumer and commercial customer payments and transactions in an agentic AI world. "Agentic AI" in payments and transactions refer to AI systems that can take autonomous actions and make independent decisions to achieve specific user goals within a digital commerce environment (such as "find these Nike sneaker variants up to $250 in price, purchase them, and arrange for home delivery").

An increasing share of commerce and payment activity will be initiated by software agents outside of pilots, as shared protocols, governance models, and accountability frameworks compete for adoption across the value chain. AI-driven traffic to US retail websites increased 4,700 percent in 2025.[1] Tech and payments leaders are already betting on the shift to AI-driven digital commerce interfaces, and a growing wave of AI startups are also emerging, with a combination of the two developing the building blocks for fully autonomous shopping.[2] Retail AI agents are moving beyond customer support to play a larger role in personalization and shopper engagement, including payment transaction authorization within certain parameters.

Issuing financial institutions must also pay attention to the various standards emerging from payment networks (and any future standards). Each current approach below places different emphasis on identity, intent, payment control, and standard setting:

- **Visa Trusted Agent Protocol (TAP):** Visa is emphasizing identify verification by verifying the "who" behind the AI agent. Visa's TAP is tied to Visa's card network and seeks to cryptographically verify in real time that an AI agent making a purchase is indeed legitimate and truly acting on the purchaser's behalf.[3]
- **Mastercard Agent Pay:** Mastercard is emphasizing tokenization, restricting the "how" of the agentic AI transaction. Mastercard Agent Pay builds on Mastercard's existing tokenization capabilities, creating "Mastercard Agentic Tokens." Mastercard is also partnering with Microsoft Azure OpenAI Service and Copilot Studio to establish pathways for AI systems to complete purchases within conversational interfaces.[4]
- **Google Agent Payment Protocol (AP2):** Google is emphasizing intent mandates by being able to cryptographically prove the "what" and "why." AP2 is an open, payment agnostic standard for agents to transact via cards, bank transfers, or even stablecoins and cryptocurrency, using cryptographic user mandates to prove consent.[5]
- **Stripe and OpenAI Agentic Commerce Protocol (ACP):** Stripe and OpenAI are emphasizing standardized discovery and structuring the "where" to reduce friction and ambiguity by using standard setting and discoverability. ACP is an open-source solution focused on "conversational" checkout and seamless purchase and utilizes shared payment tokens for AI-mediated transactions in chats/apps.[6]

Fraud Shifts to Agentic AI/Agent Manipulation

Fraud will increasingly target agent-driven workflows rather than individual accounts or cards. Attackers will influence outcomes through input manipulation, synthetic interactions, and falsified context. Any issuing bank familiar with the current state of digital commerce knows the landscape of federal and state regulations and statutes, case law rulings, and payment network rules that set the framework under which a merchant must prove the purchaser's intent and authorization to make a transaction.[7]

Where agentic AI adds a wrinkle to the current framework is as follows:

- **Current State:** Under current checkout and payment flows, the human/company making the purchase is involved in both the Point of Intent ("I want to buy this") and the Point of Checkout ("I authorize the purchase with my credit card").

[1] Deep Dive: The Role of Visa's Trusted Agent Protocol in Agentic Commerce, Sam Boboev, *Fintech Wrap Up*, October 19, 2025.

[2] 3 markets fueling the shift to agentic commerce, CB Insights, August 4, 2025.

[3] Deep Dive: The Role of Visa's Trusted Agent Protocol in Agentic Commerce, Sam Boboev, *Fintech Wrap Up*, October 19, 2025.

[4] Mastercard Launches Agent Pay for AI Payment Transactions, Louis Thompsett, *Fintech Magazine*, May 2, 2025.

[5] Google Launches New Protocol for Agent-Driven Purchases, Russell Brandon, TechCrunch, September 16, 2025.

[6] How OpenAI and Stripe's Latest Move Could Blow Up Online Shopping As We Know It, Sharon Goldman, *Fortune*, September 20, 2025.

[7] While too long for this article, such existing digital commerce legal framework includes: (a) federal and state statutes including the Federal Electronic Signatures in Global Commerce (E-SIGN) Act and state versions of the Uniform Electronic Transactions Act (UETA) portion of the Uniform Commercial Code (UCC) (except for New York, which has its own "Electronic Signature Records Act" (NY State Tech. Law §301 *et seq.*)); (b) case law rulings holding the enforceability of "shrinkwrap"/"clickwrap" terms of use agreements; and (c) payment network rules, include requirements from private payment networks such as Nacha (for ACH transactions), Visa, Mastercard, American Express, and Discover) regarding required end user/cardholder transaction authorization and retention requirements.

- **Future State Under Agentic AI:** Under agentic AI checkout and payment flows, the Point of Intent and the Point of Checkout are separated for the first time:
    - » The Point of Intent stays with the human who is delegating to the AI agent, and any related merchant terms and conditions probably need to stay with the human at the Point of Intent level as well as in order to be enforceable. There should never be "autonomous code" acting solely as "buyer;" rather, the authorization point should be moved up the transaction chain to where the human authorizes the AI agent to take certain actions on the human's behalf within a set of delegated parameters.
    - » The Point of Checkout is being delegated by the human to the AI agent under a set of parameters.

But the truly open question and unique issue for agentic AI transactions is who is liable when the *AI agent itself malfunctions*, such as hallucinating a transaction the human user did not authorize, or exceeding the boundaries of the authority delegated to it (e.g., buying 25 pairs of sneakers instead of 2, as instructed by the human user). The company developing the AI agent may try to disclaim all liability, along with direct and indirect damages in its terms of use. But if that is allowed, who gets stuck with the erroneous transaction loss "hot potato"—the user, the merchant, or the issuing bank for the payment method? Financial institutions, especially issuers, need to understand this emerging liability scenario with regard to any proposed agentic AI frameworks that the financial institution will have to work with to investigate alleged fraudulent or erroneous payment transactions.

Rise of AI-Native Fintechs

A new generation of fintechs is being built with AI embedded into core operations by default, allowing them to operate at lower marginal cost and handle higher volumes versus legacy operating models. As these new fintechs launch and become potential vendors to, customers of, or even partners of financial institutions, how to diligence, monitor, and oversee such AI-native fintechs are an emerging challenge for financial institutions.

Fintech Play Moving from Breadth to Depth

Fintech competition will also be shifting from broad coverage to execution within specific industries. Potential market advantages may come from developing novel methods to handle sector-specific cash flows, risk, and workflows, favoring embedded payments/finance vertical players over horizontal platforms, and many of these efforts will incorporate AI into the products and services. Certain fintechs may choose to become experts in areas such as construction logistics, healthcare receivables, or restaurant supply chains, not just in payments in general. Such niche-focused fintechs will seek to map industry-specific pain points to financial workflows and funds flow better than the current more generic incumbents (including financial institutions). Financial institutions will need a basic understanding of these niche-focused fintechs and may even have a potential role to play as financial institution partners to such entities.

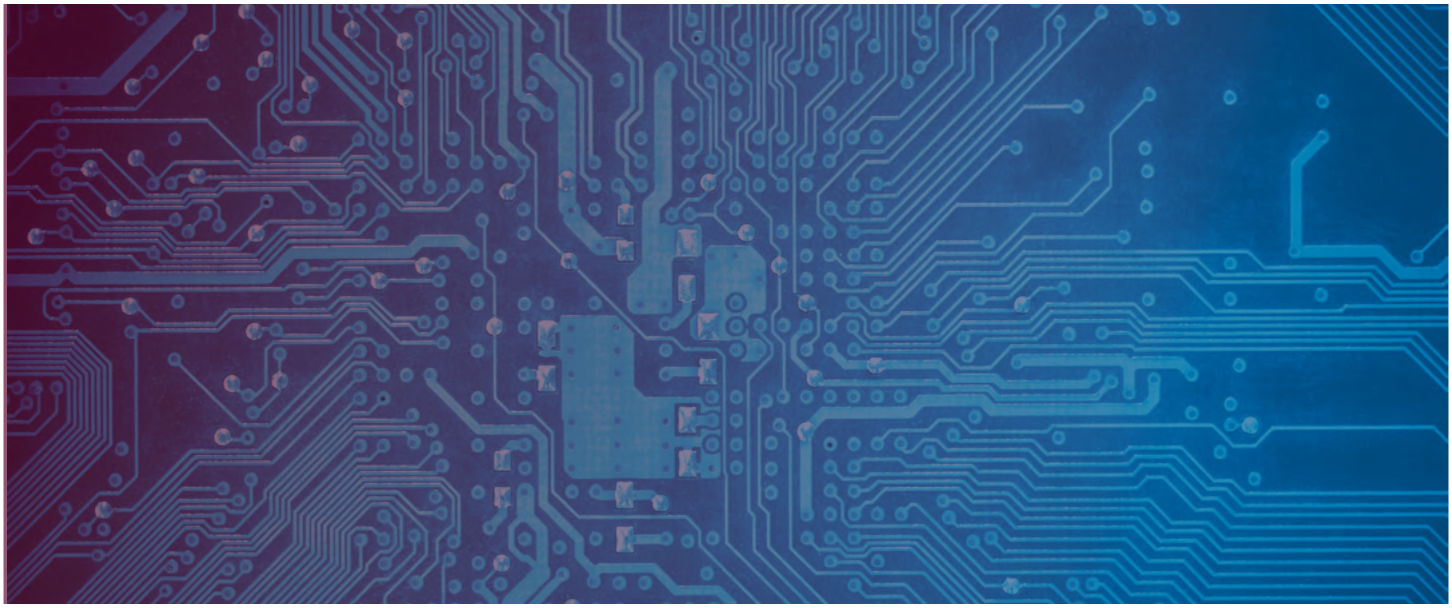

HUNTON

Hunton.com

These materials have been prepared for informational purposes only and are not legal advice. This information is not intended to create (and receipt of it does not constitute) an attorney-client or similar relationship. Please do not send us confidential information. Past successes cannot be an assurance of future success. Whether you need legal services and which lawyer you select are important decisions that should not be based solely upon these materials. Photographs are for dramatization purposes only and may include models. Likenesses do not necessarily imply current client, partnership or employee status. Contact: Samuel A. Danon, Managing Partner, Hunton Andrews Kurth LLP, 2200 Pennsylvania Avenue, NW, Washington, DC, 202.955.1500. Case results depend upon a variety of factors unique to each case. Case results do not guarantee or predict a similar result in any future case. Unless otherwise noted, attorneys not certified by the Texas Board of Legal Specialization. | 26045_01.26