The Next Move

Regulatory and policy developments in tech — June 2024

State AI laws proliferate, altering the regulatory landscape

By Jocelyn Aqua and Rohan Sen

The issue

Colorado joins a growing number of states racing to install guardrails around artificial intelligence (AI), filling the federal policy void. As a result, the US regulatory landscape for AI use looks like a patchwork of largely state-level requirements, many of them targeting specific aspects of AI use, as we await broader federal legislation.

Colorado's recently passed SB24-205, signed by the governor on May 17, 2024, is the most thorough state AI law to date. Effective on February 1, 2026, the law regulates AI use across all sectors. It imposes obligations on developers and deployers of high-risk AI systems, including enhanced duty of care, risk management and disclosure requirements, and creates various consumer rights.

Many other states — including California, Illinois, Maryland, Tennessee, Utah and Washington — have enacted bespoke legislation to regulate certain aspects of AI use. These include requirements focused on data protection and privacy, disclosure, algorithmic bias, deepfakes and government use.

Taken together, this fractured approach to AI policy could lead to overlapping requirements and compliance challenges for organizations operating in multiple states. As states continue to chart their own course, it's becoming increasingly imperative for organizations to develop a holistic governance program and compliance strategy that's broad-based yet flexible to meet most of these requirements.

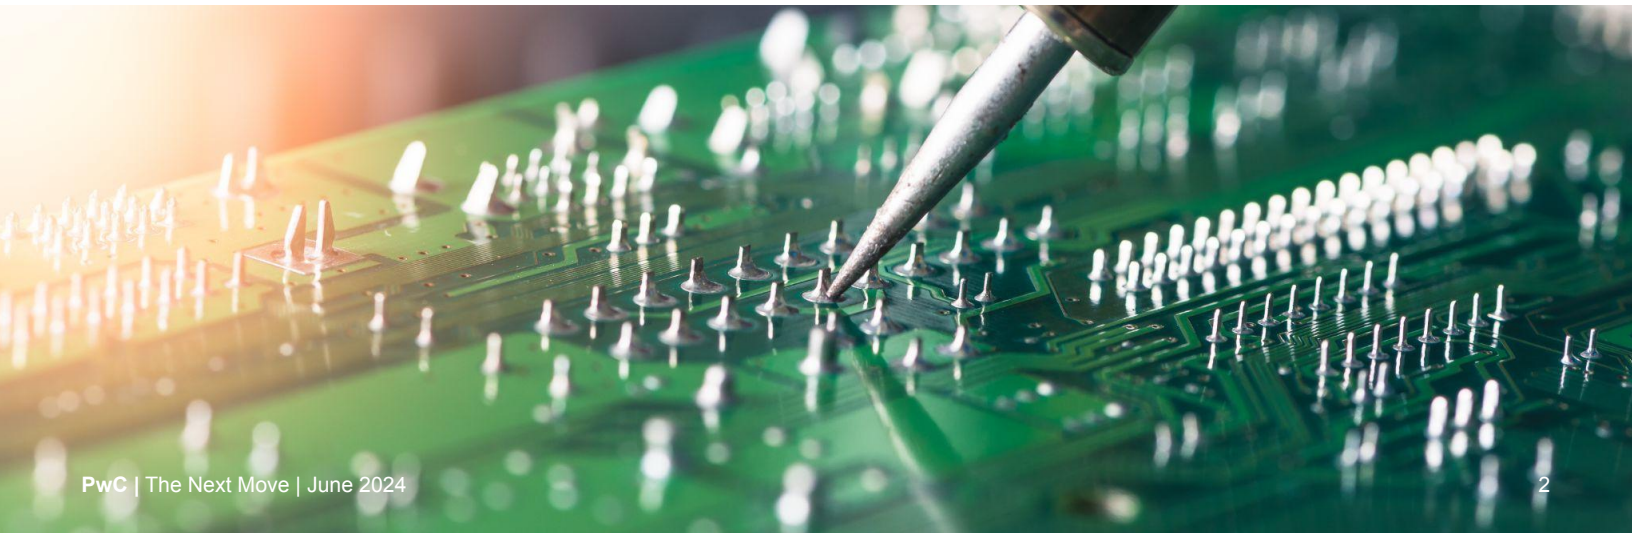

The legislatures' take

The Colorado law applies to "high-risk AI systems," defined as any AI system that, when deployed, makes or is a substantial factor in making "consequential decisions." Consequential decisions are those with a material effect on the provision, denial, cost or terms of the following to any Colorado resident:

- Education enrollment or an education opportunity
- Employment or an employment opportunity
- Financial or lending services
- Essential government services
- Healthcare services
- Housing
- Insurance
- Legal services

The law requires developers and deployers of high-risk AI systems to use reasonable care to avoid known or reasonably foreseeable risks of algorithmic discrimination.

Developer obligations. Developers must provide deployers with documentation, including high-level summaries of training data used, information on uses, risks of algorithmic discrimination and methods used to evaluate and mitigate discrimination risks. They must also publicly disclose the types of high-risk AI systems they've developed or modified and currently make available to deployers, and how the developer manages any known or reasonably foreseeable risks of algorithmic discrimination that may arise. In addition, they must disclose to the state attorney general and known deployers any known or reasonably foreseeable risk of algorithmic discrimination, within 90 days after discovery or receipt of a credible report, that the AI system has caused or is reasonably likely to have caused.

Deployer obligations. Deployers must implement a risk management policy and program, complete an impact assessment and annually review each deployment to confirm that the high-risk system isn't causing algorithmic discrimination. They must also notify consumers if the high-risk system makes a consequential decision concerning that consumer, allow them to correct errors in personal data that the system processed in making a consequential decision, and provide consumers an opportunity to appeal an adverse decision. Deployers must also disclose information about the systems they deploy and, like developers, must notify the attorney general if they discover algorithmic discrimination.

Enforcement. The Colorado attorney general has the sole authority to enforce the law and may establish rules and requirements for compliance including notice, impact assessments and developer documentation.

Other state AI laws. Legislative activity at the state level is accelerating. There are dozens of pending, enacted and failed AI bills, but certain states are leading the charge in the absence of federal legislation. Their efforts fall into several distinct categories.

The following overview isn't exhaustive but offers a glimpse into the diverse mix of enacted state laws governing AI use.

Category	Description	Enacted laws
Data protection and privacy	AI systems that rely on processing user personal data to train or function are subject to state data protection and privacy laws, which give consumers various rights such as the right to opt out of profiling processes that use sensitive personal data. Most of these laws are consumer focused, not specific to AI use.	Currently, 18 states have enacted data protection and privacy laws, including California, Colorado, Connecticut, Delaware, Iowa, Indiana, Kentucky, Maryland, Minnesota, Montana, Nebraska, New Hampshire, New Jersey, Oregon, Tennessee, Texas, Utah and Virginia.
Transparency and disclosure	Certain AI laws require disclosure so that AI use is transparent to consumers. Some require disclosing the use of bots and generative AI (GenAI) in certain contexts.	Colorado's SB24-205 requires deployers of AI systems to disclose to consumers that they're interacting with an AI system. Utah's Artificial Intelligence Policy Act requires disclosure when customers are interacting with GenAI systems or chatbots. California's BOT Act requires disclosure when bots are used to interact online for sales or political influence.
Algorithmic bias	These AI laws are aimed at preventing bias and discrimination that may result from AI systems, especially in decision-making processes that affect individual rights and opportunities.	Colorado's SB24-205 requires developers of high-risk AI systems to disclose known or foreseeable risks of algorithmic discrimination within 90 days of discovery, and its SB21-169 prohibits unfair discrimination in the context of insurance underwriting, rating and claims. In the employment context, New York City's Local Law 144 requires employers to conduct bias audits of AI tools used for employment decisions. Maryland's HB1202 prohibits employers from using facial recognition tools during interviews. Illinois' AI Video Interview Act requires employers using AI-enabled assessments to notify applicants, obtain consent and submit annual demographics reports.

Category	Description	Enacted laws
Deepfakes and synthetic content	These laws focus on combating disinformation through synthetic media or deepfakes — images, videos or voices that have been manipulated to falsely depict a person's conduct or statements.	Minnesota's HF1370 criminalizes use of deepfake technology to influence an election, as well as the dissemination of nonconsensual deepfakes depicting "intimate parts or sexual acts." Tennessee's ELVIS Act safeguards songwriters, performers and music industry professionals against commercial exploitation of their voice, image or likeness using GenAI. Alabama's HB172 prohibits a person from distributing "materially deceptive media" within 90 days before an election if they know the depiction is false and intend to harm the candidate's reputation, deceive voters and change votes. Florida's HB919 requires political ads with GenAI content falsely depicting a real person's actions with intent to injure a candidate to include a disclaimer.
Government AI use and task forces	Other laws focus on AI use by state government and/or on establishing advisory bodies to study AI use and make policy recommendations.	Texas' HB2060 establishes an AI advisory council to study and monitor AI systems developed, used or procured by state agencies. Virginia's Executive Order 30 sets standards for AI use by state agencies. Washington's SB5838 establishes a task force to assess AI use by private and public sector entities and recommend standards to regulate AI use.

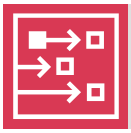

Your next move

Navigating the uneven, shifting terrain of state AI laws will require a strategic approach. By taking the following steps, businesses can help better manage the risks associated with a diverse regulatory environment and position themselves as leaders in the responsible use of AI.

Consider the following actions as you ready your organization to comply with AI requirements across multiple jurisdictions.

1. **Assess your potential exposure.** Review existing and potential state AI requirements affecting your strategy, operations, product design and compliance programs to get a preliminary view on the mitigation lift. Create a matrix that maps these requirements to your existing programs and processes, and identify gaps.
2. **Develop a compliance strategy.** Based on your potential exposure, create a plan for adapting your compliance program accordingly. Identify concrete workstreams and overlaps with other compliance obligations. Existing programs and processes can sometimes be expanded to include AI-specific measures, such as risk management, data management or cybersecurity. Consider a “solve once and for all” strategy that meets the most stringent requirements, weigh the implications (e.g., slower pace of innovation, lost business opportunity) and decide whether to take that approach or develop a bespoke solution for specific jurisdictions.
3. **Develop or enhance your AI governance model** and integrate it with your broader enterprise risk management (ERM). A critical and foundational step to developing a governance model is aligning the roles and responsibilities of existing teams, and defining new ones, to support oversight.
4. **Prepare for transparency.** If your organization faces new AI disclosure obligations, document your processes and controls, and assess their readiness for external reporting. Make sure your public statements and internal practices are aligned to stand up to increasing scrutiny from regulators, customers and the media.
5. **Monitor and adapt to evolving standards.** Track emerging requirements. Regularly conduct scenario planning exercises to prepare for possible future changes in AI regulations. This can help you quickly adapt to new laws and maintain operational continuity. Design with Responsible AI principles in mind, as that sets the baseline for your ability to be responsive to these state requirements.

FCC unveils cyber trustworthiness labels for consumer products

By Shawn Lonergan

The issue

The Federal Communications Commission (FCC) introduced a voluntary cybersecurity labeling program that establishes baseline cyber standards for evaluating wireless consumer-facing Internet of Things (IoT) products. IoT products that meet the specific criteria will qualify for the US Cyber Trust Mark, a distinctive mark analogous to Energy Star labels intended to help consumers make informed purchases, distinguish trustworthy products in the marketplace and incentivize manufacturers to develop cybersecure-by-design IoT products.

Although voluntary, the program represents an important step toward improving security for smart products and aligns with national and global efforts to safeguard consumers and promote resilience. In addition, the FCC is seeking public comment on further proposed disclosure requirements, including whether software or firmware for a product is developed or deployed in foreign adversary countries and whether customer data collected by the product will be sent to servers located in such a country.

As the federal government's first attempt to regulate the cybersecurity of consumer IoT products, the program criteria could eventually become the baseline for future cybersecurity requirements. Manufacturers should take steps to participate and get ahead of both regulatory and consumer expectations.

The regulators' take

IoT products have become increasingly essential to everyday life, but they're susceptible to many vulnerabilities. According to research cited by the FCC, cyber attacks on IoT devices surpassed 1.5 billion in the first half of 2021, and over 25 billion IoT products are expected to be in use by the end of 2030. Against this backdrop, improving IoT cybersecurity is critical.

Products in scope. The US Cyber Trust Mark program is initially open to manufacturers of wireless consumer IT products — connected or smart devices like home security cameras, network routers, garage door openers, baby monitors, TVs, thermostats, voice-activated devices, fitness trackers and more — but it may expand in the future. For now, it covers "IoT products," or IoT devices plus any product components needed to use them beyond basic operational features, like mobile apps.

Manufacturers can submit eligible products, which will be tested by accredited test labs to determine if they comply with the criteria set forth in NIST Report 8425. Adopted at the direction of Executive Order 14028 after a multiyear deliberative process involving industry stakeholders, these criteria represent the baseline capabilities that consumers can expect from IoT products.

Not eligible are medical devices regulated by the Food and Drug Administration, and equipment, devices or products from foreign entities deemed dangerous to US national security on the FCC's covered list or other lists maintained by other federal agencies, like the Department of Commerce's entity list.

The FCC's proposal seeks comment on additional national security declarations for the IoT labeling program. Ultimately, these exclusions could affect the cost of reputable products and represent a significant disruption for manufacturers with suppliers in these categories.

Labeling requirements. IoT products covered by the program and meeting its cybersecurity standards would bear a US Cyber Trust Mark label — think the Energy Star label, which indicates whether a product is energy efficient. Alongside the logo, a QR code will take users to a product registry with additional information like where to find software patches and security updates.

For wireless consumer IoT products to qualify, they must meet certain technical requirements for asset identification, product configuration, data protection, interface access control, software updates and cybersecurity state awareness. The IoT product developers must also meet requirements for documentation, information and query reception, information dissemination, and product education and awareness.

Eventually renewal will be required to keep the Cyber Trust label. When the renewal process occurs will likely depend on the type of IoT product.

Third-party collaboration. Participating manufacturers must have their product tested by an accredited and lead administrator-recognized lab and obtain product label certification by a cybersecurity label administrator (CLA). While the FCC will oversee the labeling program, a lead administrator will act as a facilitator between the agency and CLAs and be responsible for stakeholder outreach, complaint management and approval of the labs authorized to perform testing, among other things.

Litigation considerations. Because the FCC rejected industry's call for a liability safe harbor, prospective participants should consider how the label could be used in litigation. For instance, materials submitted to the labeling program could be subject to discovery in consumer protection actions if the IoT product experienced a security incident or caused injury.

Parallel cyber resilience efforts. The FCC's program is one of several moves to improve US cybersecurity under the Biden administration's national cybersecurity strategy and aligns with similar programs in Europe and Asia. Singapore already has a cybersecurity labeling scheme, Japan recently announced its own IoT labeling program, and the EU signed on to the joint CyberSafe products action plan earlier this year.

The EU's Cyber Resilience Act, discussed in our January 2024 edition, will impose cybersecurity requirements on manufacturers and distributors of connected products such as home cameras, appliances, smart watches, toys and routers. The measure will also require demonstrating conformity and affixing a cybersecurity "seal of approval" on products before they go to market.

These various, competing frameworks underscore the importance — regardless of whether you participate in the US Cyber Trust program — of building the underlying infrastructure and capabilities needed to demonstrate cyber resilience.

Your next move

The FCC's labeling program presents an opportunity for manufacturers to differentiate their products in a marketplace where consumer demands around cybersecurity are changing. By taking steps to comply, organizations can position themselves as leaders in IoT security and gain a competitive edge.

Consider taking the following steps:

1. **Assess compliance with the program's standards.** Conduct a thorough review of your IoT products to confirm that they meet the technical requirements outlined by NIST Report 8425. If you're selling products into the European Union, consider harmonizing your cybersecurity design to meet those standards as well.
2. **Evaluate your third-party risk.** Determine whether any of your suppliers or partners are on the covered lists maintained by federal agencies and could be excluded from the program. This can also help you mitigate the risk of sourcing components or services from entities that may compromise the security of your IoT products.
3. **Engage with CLAs.** Collaborate with CLAs to navigate the certification process and seek guidance on testing procedures, documentation requirements, and compliance criteria to help streamline the application process. Leveraging the knowledge of CLAs may help you address any compliance issues before they cause problems.
4. **Weigh the long-term implications.** Don't be fooled by the voluntary nature of the program, as it aligns with a wider regulatory push for cyber resilience in connected products both domestically and abroad. Evaluate the strategic benefits of earning the Cyber Trust Mark and its impact on brand reputation and market competitiveness. Balance the short-term costs of compliance with the potential long-term benefits of demonstrating a commitment to cybersecurity and consumer trust.

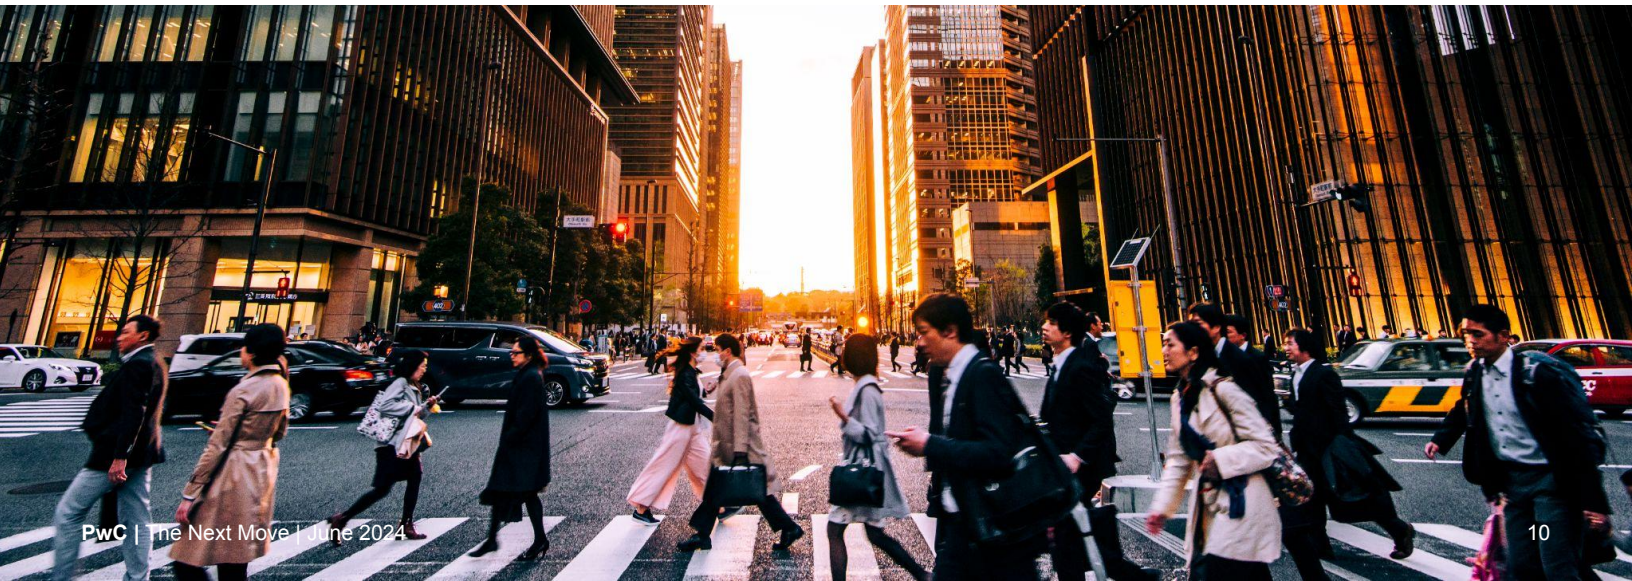

Treasury warns firms of AI risks, urges sector-wide response

By Jocelyn Aqua and Christopher Duffy

The issue

The US Treasury Department recently issued a report on the unique cybersecurity and fraud risks posed by artificial intelligence (AI) in the financial services sector. The report, mandated by President Biden's October 2023 executive order on AI, describes current use cases for cybersecurity and fraud prevention, assesses the risks presented by AI-powered threats, and reviews leading practices and recommendations for AI use by financial institutions.

The number and severity of cyber and fraud incidents continue to mount each year. According to research cited by Treasury, the average cost of a data breach reached an all-time high of $4.45 million in 2023. Online payment fraud is expected to surpass $362 billion by 2028. Losses from business email compromises exceeded $50 billion at the end of 2022. Synthetic identity fraud, which involves fraudsters leveraging the personally identifiable information of individuals, reportedly costs financial institutions more than $6 billion annually.

These concerning trends could accelerate as AI advances lower the barrier to entry for attackers, increase the sophistication and automation of attacks, and decrease time-to-exploit. Generative AI (GenAI) can help skilled threat actors develop and pilot more sophisticated malware in shorter periods of time, while helping less skilled criminals develop simple but effective attacks.

In short, AI advances have introduced new vulnerabilities and challenges for the financial sector, while also offering firms the potential for stronger defenses against these threats. Firms should heed Treasury's warnings and develop a plan to address their AI risk exposure, including by responsibly integrating the technology into their cyber and fraud prevention programs.

The regulators' take

In producing this report, Treasury conducted 42 interviews with representatives from financial institutions of all sizes and market positions, industry associations, cybersecurity and antifraud service providers, tech service providers and other stakeholders. While focused on cybersecurity and fraud, the agency also recognizes that AI use in financial services has important implications beyond these topics and will continue to study these implications.

Defensive use cases. The report notes that financial institutions have been using AI-powered fraud detection tools for more than a decade, but that recent advances have led many firms to either incorporate AI into existing threat detection tools or adopt new AI-based systems outright. "AI-driven tools are replacing or augmenting the legacy, signature-based threat detection cybersecurity approach of many financial institutions," the report found, offering the "potential to significantly improve the quality and cost efficiencies of their cybersecurity and anti-fraud management functions."

At the same time, resource requirements of AI systems may cause firms to rely increasingly on third-party IT infrastructure and data. As a result, firms "should appropriately consider how to assess and manage the risks of an extended supply chain, including potentially heightened risks with data and data processing of a wide array of vendors, data brokers, and infrastructure providers."

AI-enabled threats. The report also examines threat actors' AI use to carry out targeted cyber attacks against financial institutions. The report details four primary ways threat actors can use AI against firms with sensitive data, financial or otherwise.

- Social engineering: Using GenAI to facilitate targeted phishing, business email compromise and other fraud by enhancing culturally or locationally specific content to entice and mislead more effectively.
- Malware code generation: Using GenAI to accelerate the development of malware code, for example, to create a fake copy of a firm's website that harvests customers' credentials.
- Vulnerability discovery: Accessing AI-based tools designed for cyber defense to discover and exploit vulnerabilities in a firm's IT network.
- Disinformation: Increasing an attack's efficiency by conducting parallel disinformation campaigns using AI-generated content, such as deepfakes of company officials, to enhance their malicious campaign's influence.

Moreover, the report notes that in deploying AI-powered tools, firms are opening themselves up to unique security risks at any stage of the AI development and supply chain. These tools present novel vulnerabilities because of their dependency on the data used to train and test them, including data poisoning, data leakage, model evasion and model extraction.

Recommended actions. The report outlines priority areas for addressing AI-related operational risk, cybersecurity and fraud challenges, including:

- Improving data supply chain mapping and disclosures. This mapping would help firms understand restrictions and user rights throughout the training data supply chain and AI model-output data chain while also addressing privacy and data protection concerns. Standardized descriptions, similar to a nutrition label, for vendor-provided GenAI systems would clearly identify what data was used to train a model, where it came from, and how any data submitted to the model will be incorporated.
- Developing a common AI lexicon. AI systems use varying and imprecise terms, which can complicate efforts to identify and measure risk and to signal to users the importance of human oversight. The report includes a glossary, based on the National Institute of Standards and Technology (NIST) AI risk management framework, as a "first effort."
- Addressing capability gaps. Smaller firms generally lack the technological capabilities to create in-house AI systems and don't have the data available to adequately train systems, particularly antifraud models. The report calls for increased information sharing, creation of a data lake for fraud information, and third-party providers to develop more AI-enhanced cyber and fraud solutions.
- Expanding the NIST framework to include more substantive direction on AI governance, especially for financial firms. Treasury will assist NIST's US AI Safety Institute to establish a financial sector-specific working group charged with extending the framework toward a financial sector-specific profile.
- Developing explainability solutions. Firms that rely on third-party black box AI solutions often struggle to explain decisions to their customers and to adequately audit and test their systems.
- Creating digital identity solutions. These solutions differ in their technology, governance and security. The report supports continued research into global, industry and national digital identity technical standards.

What's next? In the coming months, Treasury will work with the private sector, other federal and state regulators, and international partners on key initiatives to address the challenges surrounding AI in the financial sector.

CFTC report. On May 2, 2024, the tech advisory committee of the Commodity Futures Trading Commission (CFTC) issued its own report on AI in the financial markets. The report makes five recommendations as to how the CFTC should approach AI, including working with industry on AI use cases and leading practices, coordinating with NIST to adapt its guidelines and framework to the financial sector, and aligning with the Securities and Exchange Commission, Treasury and other federal agencies on AI policy.

Your next move

Although the Treasury report doesn't issue any new requirements, it recognizes the increased cyber and fraud threats that AI presents and the need to use AI to better detect, prevent and mitigate those threats. These technologies are relatively nascent, and firms may not feel pressure to use them now, especially smaller firms with limited technological capabilities and access to data. However, regulators could in the future expect that firms adopt AI as part of a risk-based financial crimes program — especially as threat actors increasingly use AI.

To prepare your business for these heightened risks and expectations, consider the following steps:

1. **Inventory your AI use.** Identify all AI projects in your organization and their status (e.g., planning, development or operation). Also, identify all AI capabilities of your operational software and tools, developed internally and by third parties. This inventory will form the basis for all further decisions in establishing AI governance.
2. **Develop or enhance your AI governance model** to promote Responsible AI practices, including by third parties. Integrate this model with your broader enterprise risk management. A critical and foundational step to developing a governance model is aligning the roles and responsibilities of existing teams, and defining new ones, to support oversight.
3. **Upgrade your defenses.** Determine where AI can augment your fraud and cyber prevention capabilities. For examples, see "Keeping an even keel in shifting tides: Preventing fraud while avoiding unintended consequences."
4. **Prioritize interpretability and explainability.** Consider developing or adopting AI tools that produce outputs your employees can interpret, test, explain and audit.
5. **Mandate human oversight.** Reinforce the imperative of subjecting AI outputs and automated functions to human supervision, supported by proper training.
6. **Implement transparency.** Clearly disclose when AI is being used in customer-facing functions and allow customers to opt out.

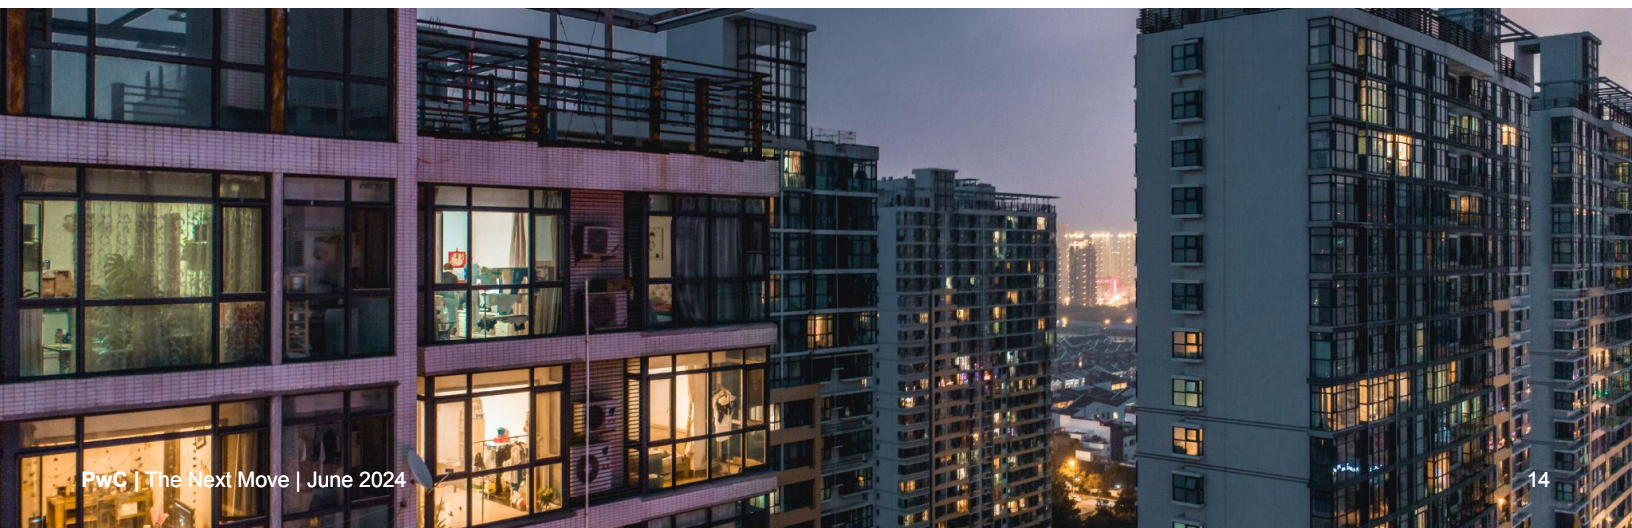

About | Contact us | Contributors

Why do we publish The Next Move?

Regulators and policymakers — keen to build new guardrails for a digital society — stand on largely unfamiliar ground. They often take different, sometimes contradictory, approaches because they have different missions and visions. At the global level, regulatory divergences reflect profoundly different value systems. Building trust in technology is complex work.

Through PwC's Next Move series, we can provide context to policy and regulatory developments in technology and tell you how you can get ahead of what might come next.

For additional information on our Next Move series, please contact:

Matt Gorham

Cyber & Privacy Innovation Institute Leader

202 951 0439

matt.gorham@pwc.com

LinkedIn

Chris Pullano

Financial Services Advisory Partner

917 520 4447

christopher.pullano@pwc.com

LinkedIn

Contributing editor and author: Ted Trautmann