

The Next Move

Regulatory and policy developments in tech — September 2023

AI companies' White House pledge: Moving from commitment to action

By [Rohan Sen](#), [Tim Persons](#), [Ilana Golbin Blumenfeld](#), [Jocelyn Aqua](#) and [Ege Gurdeniz](#)

2

Outbound investment review looms large for US investors and companies

By [Eric Lorber](#), [George Prokop](#), [Michelle Khodorov](#) and [Alison Gentry](#)

6

SEC requires prompt, robust cyber incident disclosure

By [Matt Gorham](#), [Joe Nocera](#), [Mark Cornish](#) and [Joe Sousa](#)

10



AI companies' White House pledge: Moving from commitment to action

By [Rohan Sen](#), [Tim Persons](#), [Ilana Golbin Blumenfeld](#), [Jocelyn Aqua](#) and [Ege Gurdeniz](#)



The issue

Preventing harm caused by generative AI (GenAI) to people and society is at the heart of a voluntary pledge that seven major developers of large language models (LLMs) recently [signed](#) with the US government, agreeing to place guardrails around the technology's capabilities.

The Ensuring Safe, Secure, and Trustworthy AI agreement comes amid mounting [concerns](#) about the impact to businesses and society if secure, responsible and legal practices are not followed in GenAI development and deployment. Regulators worldwide are investigating, proposing and, in some instances, beginning to adopt [rules](#) to address these concerns.

The pledge signals the White House expectation that industry “uphold the highest standards to ensure that innovation doesn't come at the expense of Americans' rights and safety.”

All companies — not just the major GenAI developers — should pay attention. As businesses of all types license these LLMs or build their own GenAI capabilities, they will become subject to similar regulatory expectations.



The administration's take

As the Biden administration [explained](#), “these commitments, which the companies have chosen to undertake immediately, underscore three principles that must be fundamental to the future of AI — safety, security, and trust — and mark a critical step toward developing responsible AI.”

Part of a broader White House push to hold AI developers accountable, the agreement is a starting point “designed to advance a generative AI legal and policy regime.” Under its terms, the companies “intend these voluntary commitments to remain in effect until regulations covering substantially the same issues come into force.”


In all, there are eight commitments. They suggest responsibilities in several key areas of the business — both internal to these organizations as well as the consumers of these technologies. Impacted stakeholders include IT, cyber R&D, trust and safety functions, corporate social responsibility and others.

Commitment

Detail

1. Conduct internal and external red-teaming (security testing) of their AI systems before their release.	This testing, which will be carried out in part by independent experts, guards against some of the most significant sources of AI risks, such as biosecurity and cybersecurity, as well as its broader societal effects.
2. Share information across the industry and with governments, civil society and academia on managing AI risks.	This includes industry leading practices for safety, information on attempts to circumvent safeguards, and technical collaboration.
3. Invest in cybersecurity and insider-threat safeguards to protect proprietary and unreleased model weights.	Model weights are the most essential part of an AI system, and the companies agree that it is vital that the model weights be released only when intended and when security risks are considered.
4. Facilitate third-party discovery and reporting of vulnerabilities in their AI systems.	Some issues may persist even after an AI system is released and a robust reporting mechanism enables them to be found and fixed quickly.
5. Develop robust technical mechanisms to confirm that users know when content is AI-generated, such as a watermarking system.	This action allows creativity with AI to flourish but reduces the dangers of fraud and deception.
6. Publicly report their AI systems' capabilities, limitations and areas of appropriate and inappropriate use.	This report will cover both security risks and societal risks, such as the effects on fairness and bias.
7. Prioritize research on the societal risks that AI systems can pose, including on avoiding harmful bias and discrimination and protecting privacy.	The track record of AI shows the insidiousness and prevalence of these dangers, and the companies commit to rolling out AI that mitigates them.
8. Develop and deploy advanced AI systems to help address society's greatest challenges.	From cancer prevention to mitigating climate change, AI — if properly managed — can contribute to the prosperity, equality and security of all.

Does the pledge stack up against AI risk management leading practices? Delivering on the eight commitments is a good start. However, the agreement only lightly touches on several other important considerations, such as data privacy and AI governance, that are central to AI risk management.



To build on trust and transparency, NIST's [AI Risk Management Framework](#) and PwC's [Responsible AI Toolkit](#) can help companies think through the structures necessary to demonstrate how they will meet these commitments.

PwC's Responsible AI Toolkit

Strategy

Data & AI Ethics

Consider the moral implication of uses of data and AI and codify them into your organization's values.

Policy & Regulation

Anticipate and understand key public policy and regulatory trends to align compliance processes.

Control

Governance

Enable oversight of systems across the three lines of defense.

Compliance

Comply with regulation, organizational policies, and industry standards.

Risk Management

Expand transitional risk detection and mitigation practices to address risks and harms unique to AI.

Responsible Practices

Interpretability & Explainability

Enable transparent model decision-making.

Sustainability

Minimize negative environmental impact and empower people.

Robustness

Enable high performing and reliable systems.

Bias & Fairness

Define and measure fairness and test systems against standards.

Security

Enhance the cybersecurity of systems.

Privacy

Develop systems that preserve data privacy.

Safety

Design and test systems to prevent physical harm.

Core Practices

Problem Formulation

Identify the concrete problem you are solving for and whether it warrants an AI/ML solution.

Standards

Follow industry standards and best practices.

Validation

Evaluate model performance and continue to iterate on design and development to improve metrics.

Monitoring

Implement continuous monitoring to identify drift and risks.

Subsequent industry collaboration. Soon after the White House pledge, four of the signers formed the [Frontier Model Forum](#). The AI-industry group intends to help realize the emphases of the commitments on research, information sharing and applications for societal benefit. It will “draw on the technical and operational expertise of its member companies to benefit the entire AI ecosystem, such as through advancing technical evaluations and benchmarks, and developing a public library of solutions to support industry best practices and standards.” The forum will explore collaboration with groups such as the [Partnership on AI](#) and [MLCommons](#).



Your next move


While only voluntary and high-level, the agreement is an important first step. For starters, by publicly committing to it, the companies are inviting customers, employees, shareholders and activist groups to hold them accountable for any shortcomings. More broadly, the agreement is a signal of where the US government is headed and aligns with what regulators in Europe, for example, are pursuing with the EU AI Act. Although voluntary now, the next step will likely be regulation.

Organizations that license, use or develop their own AI models can use these commitments to understand what federal regulators will hold AI developers to. Demonstrating that you’re [balancing the risks](#) with the rewards of innovation can go far toward establishing trust in your company — and help differentiate you from the competition. Consider taking the following steps.

- 1. Develop a plan for operationalizing the principles**, one that is informed by the pledge, as you build GenAI-based solutions. The commitments of the seven companies can become industry norms.
- 2. Understand how the government defines responsible AI.** Examine the FTC’s recent inquiry into a popular LLM (discussed in the [August report](#)) and ask yourself how you would answer those same questions for your organization. Review the [NIST AI Risk Management Framework](#) and the White House [Blueprint for an AI Bill of Rights](#). (At PwC, we’re using [GenAI to transform our own business](#) as well as our own [Responsible AI toolkit](#).)
- 3. Develop an enterprise governance model** that considers the commitments in particular and responsible AI principles more broadly. A critical and foundational step to developing a governance model is defining and contextualizing your AI terms at the use-case level. This means developing an AI risk taxonomy that standardizes key terms and metrics necessary for accurately measuring, monitoring and mitigating AI risk.
- 4. Anticipate more oversight and be ready to provide input to regulators.** The agreement signals a need for “new laws, rules, oversight, and enforcement” in the form of more executive action, bipartisan legislation and an international code of conduct for governing AI use and development worldwide. An agile and collaborative approach is emerging in the field of GenAI regulation, which requires a major adjustment for heads of regulatory affairs.

The steps you take to implement responsible AI principles now can distinguish you from the crowd when regulation arrives. To learn more about GenAI opportunities and risks, visit our [content hub](#) and read [Managing risks of generative AI](#).

Outbound investment review looms large for US investors and companies



By [Eric Lorber](#), [George Prokop](#), [Michelle Khodorov](#) and [Alison Gentry](#)



The issue

On August 9, the Biden Administration [directed](#) the Department of Treasury, Department of Commerce and other federal agencies to enact a new outbound investment review program to monitor and potentially block new investment in sensitive economic sectors of “countries of concern.” Concurrently, Treasury [released](#) an Advanced Notice of Proposed Rulemaking (ANPRM) to provide additional information and clarity — as well as solicit feedback — on the program’s scope. The focus on outbound investment review stems from the growing concern over US capital flows that threaten US strategic and defense interests.

These measures can impact the international investment landscape. As it stands, outbound investment screening will introduce risk to a variety of firms, including those in the asset wealth management, private equity and venture capital spaces, and companies with business ties to China. Investments in technology, manufacturing and cybersecurity companies will be particularly affected.

Firms need to align their compliance programs to the new requirements and cross-regulatory expectations.





The regulator's take

The proposed outbound investment regime combines targeted prohibitions and required notifications designed to limit the ability of US firms to invest in semiconductors, quantum technology and artificial intelligence sectors in named countries of concern. While the People's Republic of China (including Hong Kong and Macau) is the only country of concern named, the order leaves open the possibility of adding other countries in the future. Covered activities include mergers and acquisitions, private equity, venture capital, greenfield, joint ventures and certain debt financing schemes. Areas of focus include:

- **Semiconductors:** The proposed semiconductor restrictions reflect a concern that development of semiconductors and microelectronics technology enables the production of advanced integrated circuits that could give countries of concern a competitive advantage in military decision making and logistics. As such, US persons will be prohibited from investing in certain technologies that enable advanced integrated circuits, the design and production of advanced integrated circuits and the installation or sale of certain supercomputers. Other investments in these areas that do not meet the criteria for blocking will require US persons to notify Treasury within 30 days after the deal closes.
- **Quantum technology:** The proposed restrictions are narrowly designed to prevent countries of concern from using quantum technology to compromise cybersecurity controls that protect sensitive military communications. The restrictions vary based on designed end use, closely mirroring the existing federal export control framework. Investments in quantum computing would be prohibited, whereas prohibitions on quantum sensors, networking and communication systems would only restrict items that are made exclusively for national security or secure communications end uses. There is no current notification requirement for quantum technologies.
- **Artificial intelligence:** The proposed restrictions are designed to prevent countries of concern from using AI systems to strengthen their defense, surveillance and robotics capabilities. Treasury is considering prohibitions on US investments in AI-enabled software designed to be exclusively or primarily used for military, intelligence or mass-surveillance end uses. US persons would have to notify Treasury of transactions in AI-enabled software designed exclusively or primarily for cybersecurity, digital forensics, robotics and surveillance technology. Notifications would be required within 30 days after closing.

There are several required steps before the program can be implemented. Of note, responses to the ANPRM are due by September 28, 2023. Following the feedback period, Treasury will issue a formal proposal that incorporates feedback and provides more detail. The proposed restrictions will likely evolve over the next year and final rules will likely not take effect until late 2024.

US firms should consider carefully tracking a number of areas to prepare for final rule implementation, including:

- **Knowledge standard:** The proposed program shifts the compliance burden to US persons, who would have to determine whether a transaction is prohibited, subject to notification or permissible without notification. This requirement relies heavily on an agreed-on knowledge standard for potentially prohibited activity (i.e., criminal or civil enforcement would only happen if the US person has the requisite knowledge). Treasury is considering adopting a similar knowledge standard definition used in current export control regulations. Under this definition, US persons would have to conduct reasonable due diligence and would be held responsible for violations if publicly available information indicates that they are undertaking a prohibited transaction. This standard will likely be further refined. However, the scope of what Treasury considers reasonable due diligence will be critical in determining the compliance burden that US companies will face under the final rule.
- **Global reach:** Treasury is seeking feedback on the definition of US person but anticipates that the restrictions will apply to US persons wherever they are located. Under this definition, US subsidiaries of foreign companies would be covered, as would any investment or firm managed by a US person. It is unclear if this definition will be narrowed in the rulemaking process.
- **Lower-risk investments:** The ANPRM proposes carving out exceptions for certain types of passive and other lower-risk investments. Treasury is considering excluding publicly traded securities, index funds, mutual funds, exchange-traded funds and committed but uncalled capital investments — as well as intracompany fund transfers from a US parent company to its subsidiary — from the prohibitions and notification requirements. The scope of the exceptions is under review and will likely evolve during the rulemaking process.





Your next move

Asset management, private equity and venture capital firms, as well as companies with business ties to China, should consult with trusted advisors and consider taking immediate steps, such as:

- Engaging with relevant compliance and legal teams to review the ANPRM to understand the scope and purpose of the restrictions.
- Considering developing a point of view to express to the US government, including through relevant business organizations, on key elements of the ANPRM.
- Conducting an exposure analysis to understand the potential scope of your firm's obligations and risks when the new restrictions come into force.

To prepare for final rule implementation, prioritize the following activities over the next year:

- Assess whether your firm's planned investments will need to be submitted for regulatory notification.
- Reinforce the alignment of compliance resources with the corporate development function to best support this new regulatory screen and assess both initial investment and ongoing monitoring needs.
- Assess the design, configuration and implementation of compliance programs to facilitate review and analysis of proposed investments, including the use of technology-enabled solutions, to identify and mitigate emerging risks.
- Develop policies and procedures to notify Treasury of covered transactions within the required timeframe.
- Establish and run investment monitoring mechanisms to confirm that investments are not indirectly being channeled in violation of any requirements.
- Implement mitigation strategies and additional controls if investments of concern are identified once outbound screening requirements become final. Consider enlisting national security, cyber and privacy specialists to assist in developing proposed risk mitigation strategies, including technology-enabled persistent monitoring solutions.
- In addition to firms directly impacted by forthcoming restrictions, others with business ties to China should conduct risk and vulnerability assessments in case China responds with similar restrictions.
- Monitor regulatory updates and stay current on additional legislation that will likely continue to evolve and expand related to this area.

SEC requires prompt, robust cyber incident disclosure

By [Matt Gorham](#), [Joe Nocera](#), [Mark Cornish](#) and [Joe Sousa](#)



The issue

The SEC released its [final rule](#) on cybersecurity risk management, strategy, governance and incident disclosure on July 26, 2023. The rule aims to protect investors and the economy as a whole from the damage that a cybersecurity breach can cause. As the number, severity and cost of attacks continue to mount, investors are demanding transparency from the companies in which they've placed their resources and trust.

The final rule streamlines many of the original, proposed disclosure requirements in response to more than 150 comment letters filed from issuers, investors and other parties. Still, disclosure can seem a daunting prospect if your company's cybersecurity program won't withstand investor scrutiny. Many companies [aren't ready](#) today to reveal their cyber capabilities to the extent required.

With this new rule, the SEC puts the onus on companies to give investors current, consistent and [decision-useful](#) information about how they manage their cyber risks. Organizations should be ready to expand their disclosures regarding their cyber risk management, strategy and governance processes before the new requirements take effect in mid-December.



The regulator's take


The final rule requires registrants to disclose material cybersecurity incidents on Form 8-K. The filing is due within four days of determining that an incident is material. A filing extension is available if the US attorney general determines that immediate disclosure would pose a substantial risk to national security or public safety.

"Cybersecurity incident" means an unauthorized occurrence — or series of related occurrences — on or conducted through the company's information systems that jeopardizes the confidentiality, integrity or availability of its information systems or any information residing therein.

The final rule also requires companies to annually report material information regarding their cybersecurity risk management, strategy and governance on Form 10-K.

Category	Required disclosure
Incident reporting	<p>Report “material” cybersecurity incidents on a Form 8-K within four business days of materiality determination.</p> <p>Describe the nature, scope and timing of the incident and the material impact or reasonably likely material impact on the registrant. To the extent required information isn’t determined or isn’t available at filing time, the 8-K should disclose this fact and be later amended when the information is determined or becomes available.</p> <p>Materiality determination should be based on federal securities law materiality, including consideration of quantitative and qualitative factors.</p>
Risk management and strategy	<p>Describe the company’s process, if any, for assessing, identifying and managing material risks from cybersecurity threats, including whether the company:</p> <ul style="list-style-type: none"> • Integrated the processes into its overall risk management program • Engages consultants, auditors or other third parties, and • Has processes to oversee and identify risks from use of third-parties. <p>Describe whether and how any risks from cybersecurity threats have materially affected or are reasonably likely to materially affect the registrant’s business strategy, results of operations, or financial condition.</p>
Governance	<p>Describe the company’s governance of cybersecurity risks as it relates to:</p> <ul style="list-style-type: none"> • The board’s oversight of cybersecurity risk, including identification of any board committee or subcommittee responsible for oversight and the process by which they’re informed about cyber risks. • Management’s role and expertise in assessing and managing material cybersecurity risk and implementing cybersecurity policies, procedures and strategies. • Specific disclosure of any management positions or committees responsible for assessing and managing cyber risks, including discussion of their relevant expertise.

Effective dates. The material incident disclosure requirements take effect on December 18, 2023 (smaller reporting companies have a 180-day deferral). Requirements for risk management, strategy and governance reporting take effect for all registrants for fiscal years ending on or after December 15, 2023.



Affected stakeholders. At most companies, responsibility for compliance falls to those in several primary roles, each with its own questions to address.

1. **CEO or CFO:** When I sign and certify my company's 10-K, am I confident in the integrity and completeness and accuracy of the information the company is disclosing related to the cyber risk management program? Are we prepared to make the expanded disclosures the new rule requires?
2. **Board:** Are we getting the effective, ongoing reporting we need to understand the key cyber risks and what management is doing to mitigate those risks? How do we know we're asking the right questions of the CISO and others who report to us? Are we comfortable with our own cybersecurity knowledge to effectively oversee this area?
3. **CIO/CISO and team:** Are the details of my cyber risk management program sufficient to meet the new requirements? How much do we disclose without introducing additional risk to the company? How will we help the people responsible for determining an incident's materiality make that judgment without unreasonable delay? If an incident is material, how do we confirm that required information is included in the filed 8-K within the four-day window?
4. **Legal:** How can we draft compliant disclosures without revealing confidential information about our cyber program? Which criteria should we use to determine (with the CISO and those responsible for SEC reporting) an incident's materiality? If immediate disclosure could pose a substantial risk to public safety or national security, how do we report it to federal law enforcement and confirm we are well coordinated internally? How will we be informed of any determinations and communications made to the SEC that could affect the timing for disclosure?
5. **Internal audit:** What's our role in confirming that disclosures are complete, accurate and sound?

Enforcement. Organizations that don't comply with the new rule will likely face serious consequences, as recent SEC enforcement [actions](#) suggest. The agency has levied large fines against companies for not disclosing breaches sufficiently or in a timely manner. It continues a two-pronged approach to enforcement, requiring that organizations (1) make appropriate disclosures under the rule and (2) have controls and procedures in place to escalate necessary items for determination of whether disclosures are required.








Your next move

Complying with the new rule will require coordination among security, finance, risk and legal teams, as well as key business leaders (as appropriate). So it's helpful to agree internally on how to make timely and accurate disclosures that satisfy the SEC rule.

The good news is that these capabilities are mutually-reinforcing. Improving one area will also help you improve in other areas. Consider this framework for unifying your organizational efforts.

A PwC data security framework

Capabilities needed for decision-useful reporting on your cyber strategy and practices

SEC element	Necessary capabilities			
 Cyber risk management & strategy	Cybersecurity risk management program	Enterprise cyber risk assessment	Cybersecurity policies and standards	Monitoring & reporting of cyber resilience and posture
 Cyber incident reporting	Security event monitoring & detection (team, process, tools)	Incident and crisis response (team, process tooling)	Process to determine incident materiality	Incident register and practiced process to drive 4-day & periodic obligations
 Cyber governance	Board cybersecurity oversight	Incorporating cyber risk into business strategy, financial planning, and capital allocation	Cybersecurity board and management governance disclosed in regulatory filings (10-K)	

Companies should answer these questions clearly to pass muster with the SEC *and* investors.

- 1. What's our process for reporting incidents?** It's important for leaders and the board to understand the internal escalation and external reporting processes. Test the escalation process now, before an event occurs.
- 2. How can we effectively determine materiality?** Given the complexity, materiality determinations should not be the responsibility of any one person. Involve the CFO, general counsel, CISO, CIO and front-line business leaders. In assessing cyber materiality, you should consider qualitative factors such as effects on reputation, customer relationships, vendor relationships and regulatory compliance. And you should begin taking a long view of breaches and breach attempts, considering cumulative effects of related occurrences.
- 3. Have we documented our processes for determining materiality?** Documenting how you determined an incident's materiality is critical, particularly if you found it not material. If the SEC questions your conclusion, you'll need to justify with details of your processes and considerations of quantitative and qualitative factors and the basis for your decision.
- 4. How much disclosure is too much?** Recognizing the sensitivity of cybersecurity programs, the SEC's disclosures are generally more principles-based. Still, complying with the new requirements without revealing confidential information about your cybersecurity procedures and program will be an important consideration. Some companies are creating a standard template for reporting incidents to have on hand, then modifying it should an event occur.
- 5. Can we meet the four-day reporting deadline?** We see lots of confusion surrounding the four-day timeframe for disclosure. The clock starts ticking *not* when the incident occurs or is detected, but when it's determined to be "material." The rule does not impose any specific timeline between the incident and the materiality determination, but the materiality determination should be made without unreasonable delay.
- 6. Are we prepared to report related, material events?** The final rule removed the requirement to aggregate disparate non-material risks to determine if an 8-K disclosure is required. However, the final rule still requires events that are related — for example, by the same malicious actor or that exploit the same vulnerability — to be reported if found material.

7. Is our cyber risk management and strategy up to par?

- **Policies and procedures:** Are yours in line with the specifications in at least one recognized industry framework? Are they updated regularly? Does everyone in the organization know what they are, and how they're responsible for following them? Are they well enforced?
- **Risk assessment process:** Having a risk assessment process isn't enough. Is yours robust? Is it applied throughout the organization, focusing on top risks to the business? How often do you do risk assessments? Are assessment results incorporated into your enterprise cyber strategy, enterprise risk management program and capital allocations? Have you engaged a third party to assess your cybersecurity program?
- **Controls and controls monitoring:** How does your organization monitor the effectiveness of its risk mitigation activities and controls? How mature are your capabilities, as evaluated against an industry framework? How are leadership and the board informed about the effectiveness of these controls?

8. Does the board have enough information for oversight?


The final rule requires a description of the board's oversight of cybersecurity risks in the 10-K. This could include where cybersecurity risk is allocated (e.g., full board, committee) and the processes by which the board or committee is informed about the risks.

While the final rule does not require disclosure of cyber knowledge, directors should consider their comfort level in the areas that the rule specifies. As a whole, boards are responsible for understanding cyber concepts and requirements well enough to provide oversight. How will directors individually and collectively confirm that they continue to learn? Will they invite cyber specialists to meetings? Will they attend classes or other kinds of training? Will they consult with outside advisors?

For more detail, see [Making materiality judgments in cybersecurity incident reporting](#). Bookmark the [PwC site on SEC cyber disclosure rule](#) to find new content on the role that internal audit, IT/security, finance and legal teams can play.



Contact us



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

Chris Pullano

**Financial Services
Advisory Partner**

917 520 4447

christopher.pullano@pwc.com

Contributing editors: Ted Trautmann, Cristina Ampil

© 2023 PwC. All rights reserved. PwC refers to the US member firm or one of its subsidiaries or affiliates, and may sometimes refer to the PwC network. Each member firm is a separate legal entity. Please see www.pwc.com/structure for further details.

This content is for general information purposes only, and should not be used as a substitute for consultation with professional advisors. PwC US helps organizations and individuals create the value they're looking for. We're a member of the PwC network of firms in 155 countries with more than 327,000 people. We're committed to delivering quality in assurance, tax, and advisory services. Tell us what matters to you and find out more by visiting us at www.pwc.com/us 892038-2021 AP CT