



MAYER | BROWN

# SECURING AGENTIC AI

## OPPORTUNITIES FOR LEGAL TEAMS TO REDUCE RISK

October 15, 2026

## TODAY'S SPEAKERS



PARTNER

**VERONICA GLICK**

CHICAGO

VGLICK@MAYERBROWN.COM



PARTNER

**STEPHEN LILLEY**

WASHINGTON DC AND NORTHERN  
CALIFORNIA

SLILLEY@MAYERBROWN.COM

## AGENDA


1. Agentic AI: A Diverse, Transformative Technology In A Nutshell
2. Error Code: When Agentic AI Goes Wrong
3. Securing Agentic AI: Reducing Legal Risk
4. 10 Questions In-House Counsel Should Ask

The background is a complex digital-themed abstract. On the left, there are glowing orange and yellow circuit traces and patterns. On the right, there are blue and teal circuit traces and vertical columns of binary code (0s and 1s). The overall color palette is a mix of warm oranges and cool blues.

# 01

AGENTIC AI: A DIVERSE, TRANSFORMATIVE  
TECHNOLOGY IN A NUTSHELL


# Agents, agents, everywhere



**INTEGRATOR (optional overlay)** Builds and/or operates agents. The contract chain becomes Company ↔ Integrator ↔ Vendor(s).

# AGENTIC AI CHANGES THE SECURITY CALCULUS

- Agents autonomously plan and execute steps—which may include **accessing data, executing code, overriding approval workflows, or distributing information externally**—without additional human authorization.
- The broad access and autonomy associated with many agentic AI implementations can expand security risks.
- A compromise, configuration error, or misaligned instruction may create **consequences that unfold faster and more broadly** than in conventional software environments.



The background is a complex digital-themed abstract. It features glowing orange and yellow circuit traces on the left side, transitioning into a dark blue field with faint, glowing blue circuit patterns. On the right side, there are vertical columns of glowing blue binary code (0s and 1s). A thin white vertical line runs down the center of the image, passing behind the text.

# 02

ERROR CODE: WHEN AGENTIC AI GOES WRONG


## AGENTS IN THE NEWS

- Recent, widely reported incidents involved AI agents **compromising third-party systems**, impacting both public and private sector entities.
- These incidents demonstrate the power of today's models and the **challenge of keeping that power contained**.
  - Comparable capabilities are proliferating among developers, vendors, and threat actors, so these risks are not confined to frontier AI developers.
- **Responding is highly complex**: a single incident may involve thousands of automated actions across many systems and third parties.
  - Effective response may depend on the timely preservation and analysis of extensive forensic records, including agent logs, prompts, tool activity, and records held by third parties.
- Rogue agents are only one risk; the broad use of agentic tools can go wrong in ways that are serious, costly, and headline-making.




## AGENTS IN THE NEWS

- Recent, widely reported incidents involved AI agents compromising third-party systems, both public and private.
- These incidents demonstrate the power of today's models—and the challenge of keeping that power contained.
- Front-page proof that things can go wrong in significant ways.
- Not just a problem for frontier labs: as access spreads, everyone will have equivalent power—including bad actors.
- Many ways things can go wrong—rogue agents are only one:
  - Attackers using agents to scale cyberattacks
  - Manipulation through prompt injection
  - Over-permissioned or misconfigured agents
  - Destructive mistakes in live systems
  - Compromise through vendors and other third parties



# Agentic Top 10 At A Glance



**ASI01:** Agent Goal Hijack

**ASI03:** Identity & Privilege Abuse

**ASI05:** Unexpected Code Execution (RCE)

**ASI07:** Insecure Inter-Agent Communication

**ASI09:** Human-Agent Trust Exploitation

**ASI02:** Tool Misuse & Exploitation

**ASI04:** Agentic Supply Chain Vulnerabilities

**ASI06:** Memory & Context Poisoning

**ASI08:** Cascading Failures

**ASI10:** Rogue Agents

# LEADERS HAVE WARNED ABOUT SECURITY RISKS

CISA / Five Eyes — Joint guidance “Careful Adoption of Agentic AI Services” (May 1, 2026)


- Identifies key categories of risks presented by agentic AI, including:
  - Privilege risks
  - Behavior risks
  - Structural risks
- Establishes best practices for securing agentic AI systems, such as controlling context windows of agents; designing agents with oversight mechanisms; and constructing strong identity management.





# 03


SECURING AGENTIC AI: REDUCING LEGAL RISK



## HOW CAN LAWYERS REDUCE LEGAL RISK ASSOCIATED WITH THE SECURITY OF AGENTIC AI?


- Companies face potentially significant risks if AI agents are abused or act in unintended ways.
- Security models, expert guidance, and tooling to manage AI agents are emerging rapidly.
- But securing agentic AI remains fundamentally difficult, including because of:
  - Core tension between agentic capability and risk
  - Complexity of security in the agentic context, including because of prompt injection and other AI-specific attacks
  - Reliance on third parties
- Companies also can make this task harder through:
  - Prioritization of innovation over security
  - Persistent tension between security and AI teams
  - Lack of visibility into adoption and use of agents
  - Lack of standardization in system design
  - Weak underlying security fundamentals
- So, what can lawyers do?

# OPPORTUNITIES TO MITIGATE AGENTIC AI RISK



# OPPORTUNITIES IN GOVERNANCE AND DESIGN OF AGENTIC AI

- Governance
  - Minimizing scope creep in approved implementations
  - Maximizing visibility into implementations across the company, including those built internally
  - Increased scrutiny and testing of high-risk agentic implementations
  - Ensuring adequate training is provided to relevant engineering teams and other responsible parties
  - Auditing against regulatory guidance as applicable
- System design
  - Ensuring the security team believes that appropriate security controls have been implemented
  - Confirming that the activities of AI agents are adequately logged and appropriate alerting is in place
  - Confirming that system connections, interdependencies, and permissions are adequately understood



## OPPORTUNITIES TO ENHANCE PREPAREDNESS

- Updating incident response plans to account for likely scenarios may help the company respond more effectively when an incident occurs
- Tabletop exercises targeted at agentic AI incidents can help build readiness and work through some of the potentially complex dynamics that may come into play
- Key issues to attend to in these preparation activities may include:
  - The scope of relevant plans that were originally designed to address inbound security issues, not system malfunctions or outbound misconduct.
  - The allocation of core roles (e.g., investigation, leadership) across teams (e.g., legal, security, AI, IT) when an issue cuts across responsibilities in new ways
  - The adequacy of existing technical tools and telemetry to support the functions that the team will want to pursue – e.g., available logging.
  - Responsibility for critical decisions in the response process – e.g., turning off an important AI system to achieve containment.
  - Key dependencies, including the extent to which response will depend on information from third parties.
  - Extent to which an incident may require new capabilities, such as communicating to third parties about injury caused by an agent.


## OPPORTUNITIES IN CONTRACTING FOR SECURE AGENTIC AI

- To the extent that agentic AI solutions are provided or supported by third parties, companies can use contractual terms to mitigate some security risk.
- Key contractual terms are likely to focus on:
  - *Security standards & baseline controls*, addressing how compliance is demonstrated, the level of architectural transparency, and the shared responsibility allocation between provider and company
  - *Incident notification & response*, reflecting that incident reporting takes on heightened importance when agents can act before humans detect a problem.
  - *Audit rights & ongoing assurance*, recognizing that agentic AI may require specialized expertise for meaningful audits and ongoing assurance.
  - *Governance, oversight & termination*, with focus typically including boundaries on agent autonomy.
  - *Liability allocation*, accounting, to the extent possible, for harm from interacting causes that may be harder to untangle than in other disputes.

The background is a complex digital-themed abstract. On the left, there are glowing orange and yellow circuit traces and patterns. On the right, there are blue and teal circuit traces and patterns. In the center-right, there are vertical columns of binary code (0s and 1s) in a light blue color. The overall color palette is dominated by warm oranges and yellows on the left, transitioning to cool blues and teals on the right.


# 04

10 QUESTIONS IN-HOUSE COUNSEL CAN ASK



## KEY QUESTIONS FOR IN-HOUSE COUNSEL TO ASK

1. Are high-risk agentic implementations subject to meaningful scrutiny and testing?
2. Have business leaders and engineering teams received training on risk mitigation?
3. Does the security team believe that appropriate security controls have been implemented?
4. Are the activities of AI agents logged and subject to alerting?
5. If the company needed to shut down or mitigate the actions of AI agents, would it be able to do so?



## KEY QUESTIONS FOR IN-HOUSE COUNSEL TO ASK

6. If an agentic incident occurred tomorrow, would it be possible to quickly assemble the right group of stakeholders, each of whom understands their role and responsibilities?
7. Is it broadly understood who has responsibility for critical containment decisions?
8. Do we have a documented plan for how to respond to an agentic incident, whether as part of the cyber incident response plan or in another document?
9. Have we practiced responding to likely scenarios, either through a formal tabletop exercise or more informal scenario-based sessions?
10. Do our contracts mitigate risk to the extent reasonably possible given market realities?



# MAYER|BROWN

This Mayer Brown publication provides information and comments on legal issues and developments of interest to our clients and friends. The foregoing is not a comprehensive treatment of the subject matter covered and is not intended to provide legal advice. Readers should seek legal advice before taking any action with respect to the matters discussed herein.

Mayer Brown is a global legal services provider comprising associated legal practices that are separate entities, including Mayer Brown LLP (Illinois, USA), Mayer Brown International LLP (England & Wales), Mayer Brown Hong Kong LLP (a Hong Kong limited liability partnership) and Taill & Chequer Advogados (a Brazilian law partnership) (collectively, the "Mayer Brown Practices"). The Mayer Brown Practices are established in various jurisdictions and may be a legal person or a partnership. PK Wong LLC ("PKW") is the constituent Singapore law practice of our licensed joint law venture in Singapore, Mayer Brown PK Wong Pte. Ltd. More information about the individual Mayer Brown Practices and PKW can be found in the Legal Notices section of our website. "Mayer Brown" and the Mayer Brown logo are the trademarks of Mayer Brown. © 2026 Mayer Brown. All rights reserved.