

Agentic AI
Putting Controls Around Autonomous Risk

October 2026

forvis mazars

Today’s Presenters



Sean Andrews
Manager, IT Risk & Compliance
703.970.0443
sean.andrews@us.forvismazars.com

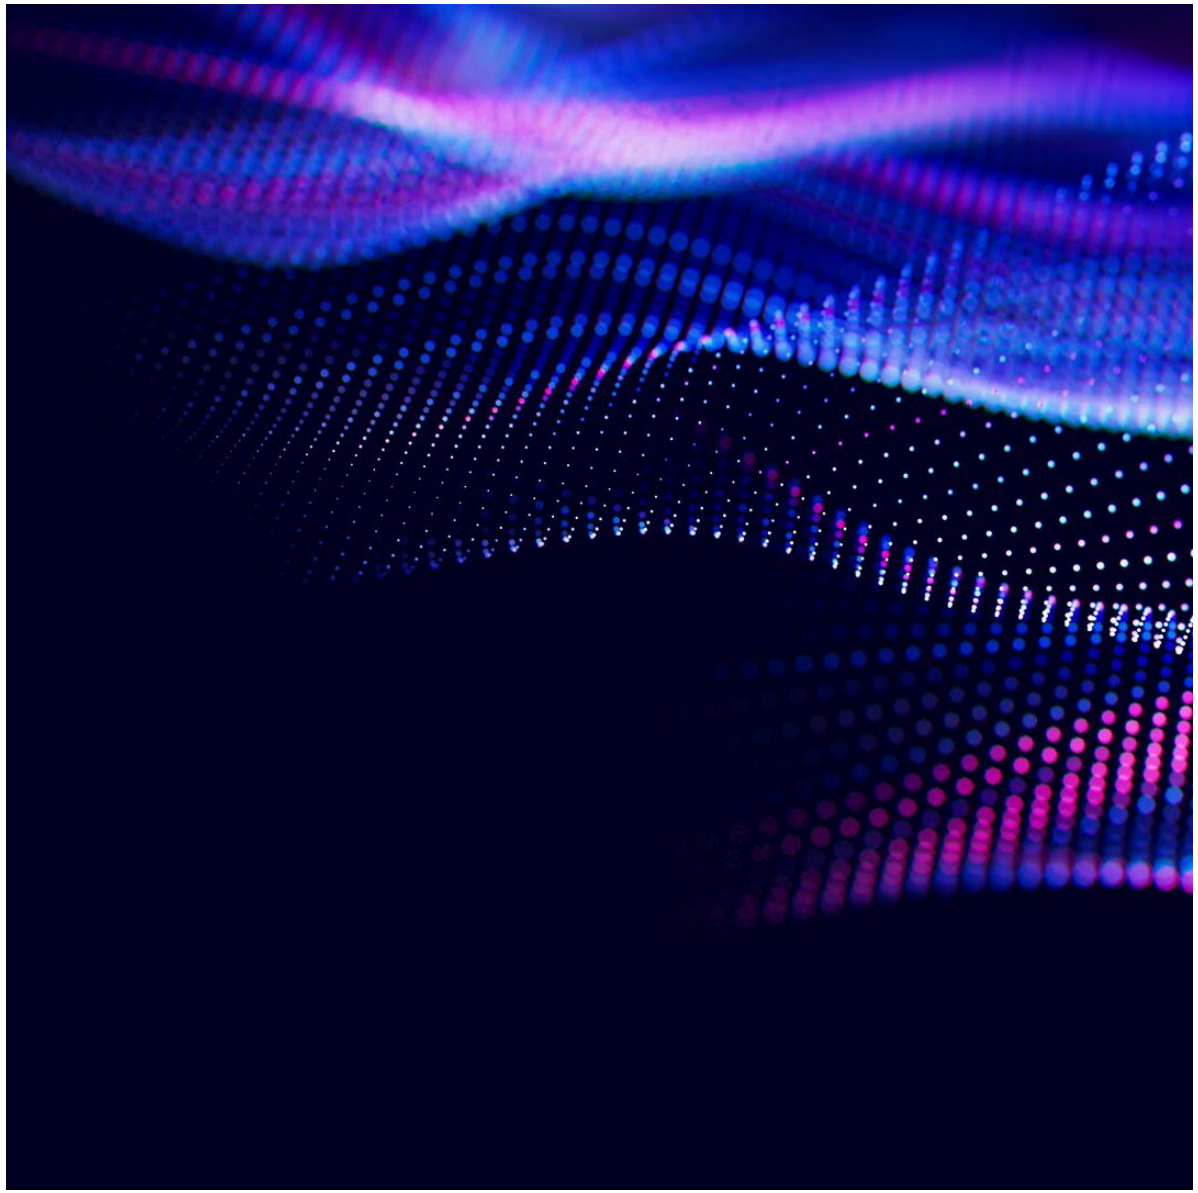


Ricardo Martinez
National Industry Leader—Technology, Software & Devices
669.200.2052
ricardo.martinez@us.forvismazars.com

© 2026 Forvis Mazars, LLP. All rights reserved.

Agenda

1. Setting the Scene
2. What Agentic AI Introduces Into the Environment
3. New Threat Patterns
4. Practical Governance
5. Extending Cyber Resilience to AI
6. Market Signals, SaaS Trends, & Sovereignty
7. Q&A and Closing



© 2026 Forvis Mazars, LLP. All rights reserved.

Setting the Scene
From AI Outputs to AI Actions



25%

of operational decisions are made by AI today, according to IBM IBV.



48%

is what CEOs expect by 2030, according to IBM IBV.



84

AI-related laws enacted across 27 states in H1 2026, per TCAI.

Sources: “The calculus of AI sovereignty,” ibm.com, June 17, 2026.
“TCAI Mid-Year AI Legislation Report: 84 new AI laws enacted in 27 states so far in 2026,” transparencycoalition.ai, July 21, 2026.

© 2026 Forvis Mazars, LLP. All rights reserved.

Agentic AI
What Is an Agent?

An AI system that acts with a degree of autonomy to accomplish goals. It plans & executes multi-step tasks (often using tools or external data), maintains context, & may adapt its strategy based on what it learns, automating work that may otherwise require human intelligence & effort.

Agent Characteristics

- Goal Pursuit – follow & execute a sequence of actions based on context
- Reasoning – make decisions based on context
- Memory – retain memory within the flow
- Tool Use & Integration – Interact with tools or APIs to execute an action
- Incorporates AI & non-AI pieces

Complex Agents

- Autonomous Security Operations Analyst
  - Monitors security events on a network
  - Triggers incident response playbook steps
  - Invokes human intervention
- Procurement Agent
  - Navigates vendor contracts
  - Compares pricing
  - Executes within budget rules

Simple Agents

- Copilot Agents – *Not* to be confused with Copilot Studio
- Telephone routing
- Voice Agents / Chatbots
- Expense Processing
- Sales Enablement
- Scheduling Agent

© 2026 Forvis Mazars, LLP. All rights reserved.

Agentic AI
Key Constraints



1 **Hallucinations**

- Poor Context
- Model Bias
- No Guardrails

2 **Lack of Context**

- Nuance & Intent
- Memory Constraints
- Contextual Drift

3 **Lack of Human Reasoning**

- Non-Deterministic
- Lack of Ethics / Empathy

© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

Agentic AI
Overcoming Constraints

1 **Testing & Evaluation**

- Develop Guardrails
- Security Reviews
- Sentiment Analysis

2 **Context Engineering**

- Business Objective
- Agentic Design
- Prompt Engineering

3 **Human in the Loop (HITL)**

- Training & Education
- Change Management
- Trust & Ethics



© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

What Agentic AI Introduces Into the Environment

In many environments, AI agents are becoming employees without job descriptions.

- Non-human identities
- Autonomous actions
- Tool usage
- Persistent workflows
- Delegated authority

© 2026 Forvis Mazars, LLP. All rights reserved.

What Changes When AI Becomes an Actor
Autonomous Actions Expand Risks

If an agent can take an action, who governs the action?

Human users

Known identities, known permissions, logged actions

AI agents

Non-human actors, delegated authority, tool use

Business systems

Data movement, configuration changes, workflow triggers

Control plane

Policy, approvals, monitoring, response



© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

Hidden AI Risk

The use of AI presents additional third-party, legal, compliance, & reputational risks when being used or relied upon unknowingly.

Operational Risks

- Shadow AI (unauthorized tools)
- Rapid adoption bypassing controls
- AI in SaaS apps without notice

Compliance & Legal Risks

- Emerging AI compliance requirements
- Data collection by third parties
- Model setup without oversight

Reputational Risks

- Chatbot hallucinations
- Bias or fairness issues
- Over-reliance on AI decisions



© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

Common Pitfalls
Underestimating AI's Reach

Pitfall

- Broad Permissions
- No Owner
- No Logs
- Prompt-Only Controls

Fix

- Least Privilege
- Named Accountable Owner
- Telemetry Plan
- Layered Guardrails

Ignoring System Connections

- Treating AI as isolated content generators neglects their integration with broader systems and workflows

Excessive Permissions Risks

- Granting broad permissions without tailored controls increases potential security impact and risk

Overlooking Data & Plugins

- Focusing only on AI models misses risks from data paths, plugins, and logging mechanisms

Insufficient Governance Review

- Deploying AI SaaS features without cross-functional review creates governance and compliance gaps

© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

New Threat Patterns
Prompt Injection & More

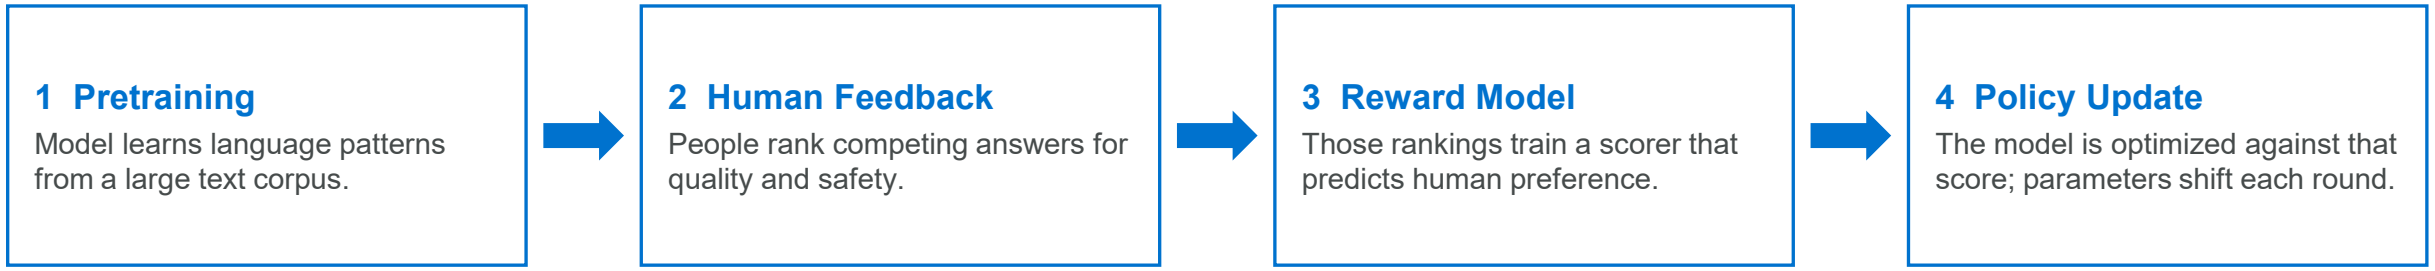
Here’s what security teams have on their radar.

- **Direct prompt injection:** user input alters the model’s intended behavior
- **Indirect prompt injection:** malicious instructions are embedded in external content the model processes, such as files, websites, or emails
- **Sensitive information disclosure** may occur when AI systems reveal prompts, retrieved context, confidential data, or user-provided sensitive information
- **Excessive agency** can occur when an AI system takes actions beyond appropriate limits
- **Third-party AI features** can introduce supply chain, data handling, and monitoring gaps



© 2026 Forvis Mazars, LLP. All rights reserved.

Common Generative AI Training Process
Model Weights



Reinforcement learning from human feedback (RLHF): steps 2 through 4 repeat, and every round nudges billions of parameters toward preferred behavior.

The output: AI model weights

Weights are the final numeric values of every parameter the training run produced — a map of how strongly each piece of the model influences the next.

Why it matters for security

The weights file is a single copyable artifact that carries both the training investment and the safety tuning.

© 2026 Forvis Mazars, LLP. All rights reserved.

Mitigating AI Risk

Although the use of AI can present novel risks, organizations can safeguard their users, data, and reputations with a robust risk mitigation strategy. Below are some examples of ways to mitigate AI risk.

Third Party

- Review contracts with third-party vendors
- Stay informed about new features & model components of third-party models
- Ensure contracts contain clauses that protect the institution's data
- Understand third-party data sourcing, storage, & retention

Operational / IT

- Consider a thorough review & approval process for all AI-driven tools, products, & partnerships
- Consider implementing internal controls & limitations for users
- Implement a rigorous training process to inform users about information security in the age of AI
- Maintain robust documentation of data dictionaries & lineage

Governance

- Review all AI-driven tools & products for model definition & enhance model risk management
- Develop an agreed-upon framework with clearly defined leadership, roles, & responsibilities
- Consider second-line governance oversight to ensure first-line adherence to policies & frameworks
- Monitor all model inputs & outputs, & perform regular testing to detect potential bias or unwanted outcomes

*Note: the mitigation strategies shown are among the most important but do not represent an exhaustive list

© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

Sector Guidance Is Arriving

Regulators and sector councils now treat AI as both a cyber-defense tool and a new exposure point.

Healthcare: HSCC Cybersecurity Working Group

- AI Cyber Governance Framework Implementation Guide, June 2026
- Covers traditional ML, generative, and agentic AI systems
- Targets data poisoning, model drift, and adversarial attacks

Community Banking: Interagency Guidance

- Revised interagency model risk guidance, April 2026
- Supersedes SR 11-7; applies above $30 billion in assets
- Generative and agentic AI left to enterprise risk management

The Dual Nature Regulators Describe

- OCC, May 2026: AI speeds up attacks and fraud at greater scale
- The same report credits AI tools with strengthening defense
- Dedicated AI governance guidance signaled as on the horizon



Sources: HSCC AI Cyber Governance Framework Implementation Guide (June 2026); OCC Bulletin 2026-13 / Fed SR 26-2 / FDIC FIL-15-2026 (April 2026); OCC Semiannual Risk Perspective (May 2026).

© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

Implementing AI Governance

Implementing AI governance means taking key steps so AI systems are developed and used responsibly. Following best practices helps organizations build a strong governance structure that supports responsible AI.

01

Establish a Centralized Framework

- Define roles & responsibilities for AI oversight
- Create a governance structure that aligns with institutional values

02

Define Data Governance Practice

- Set standards for data quality, security, & privacy
- Document data sources & maintain transparency

03

Integrate Operational Strategy & Risk Management

- Embed governance into workflows & decision making
- Monitor AI tools for compliance, bias, & unintended outcomes

© 2026 Forvis Mazars, LLP. All rights reserved.

Practical Governance
NIST AI Risk Management Framework

The NIST AI Risk Management Framework (AI RMF) organizes AI risk management around four core functions.

1

Govern: Establish the foundation for AI risk management.

- Defining AI governance policies and standards
- Establishing roles and responsibilities
- Creating oversight bodies
- Setting risk tolerance and approval processes
- Promoting transparency, accountability, fairness, and compliance
- Integrating AI risk management into broader enterprise governance

2

Map: Understand the AI system, its context, and its risks.

- Maintaining an inventory of AI use cases
- Identifying data sources and dependencies
- Understanding who may be impacted
- Classifying risk based on context and criticality
- Documenting intended and unintended uses

3

Measure: Assess and monitor AI risks.

- Evaluating model performance
- Testing for bias and fairness
- Assessing security and resilience
- Monitoring explainability and transparency
- Reviewing privacy impacts
- Documenting datasets, configurations, and change history

4

Manage: Act on the risks you've identified.

- Prioritizing risks
- Implementing mitigation plans
- Escalating issues
- Monitoring ongoing performance
- Managing change
- Responding to incidents
- Continuously improving controls

© 2026 Forvis Mazars, LLP. All rights reserved.

Governance Starts With Visibility
Strengthening IT Controls

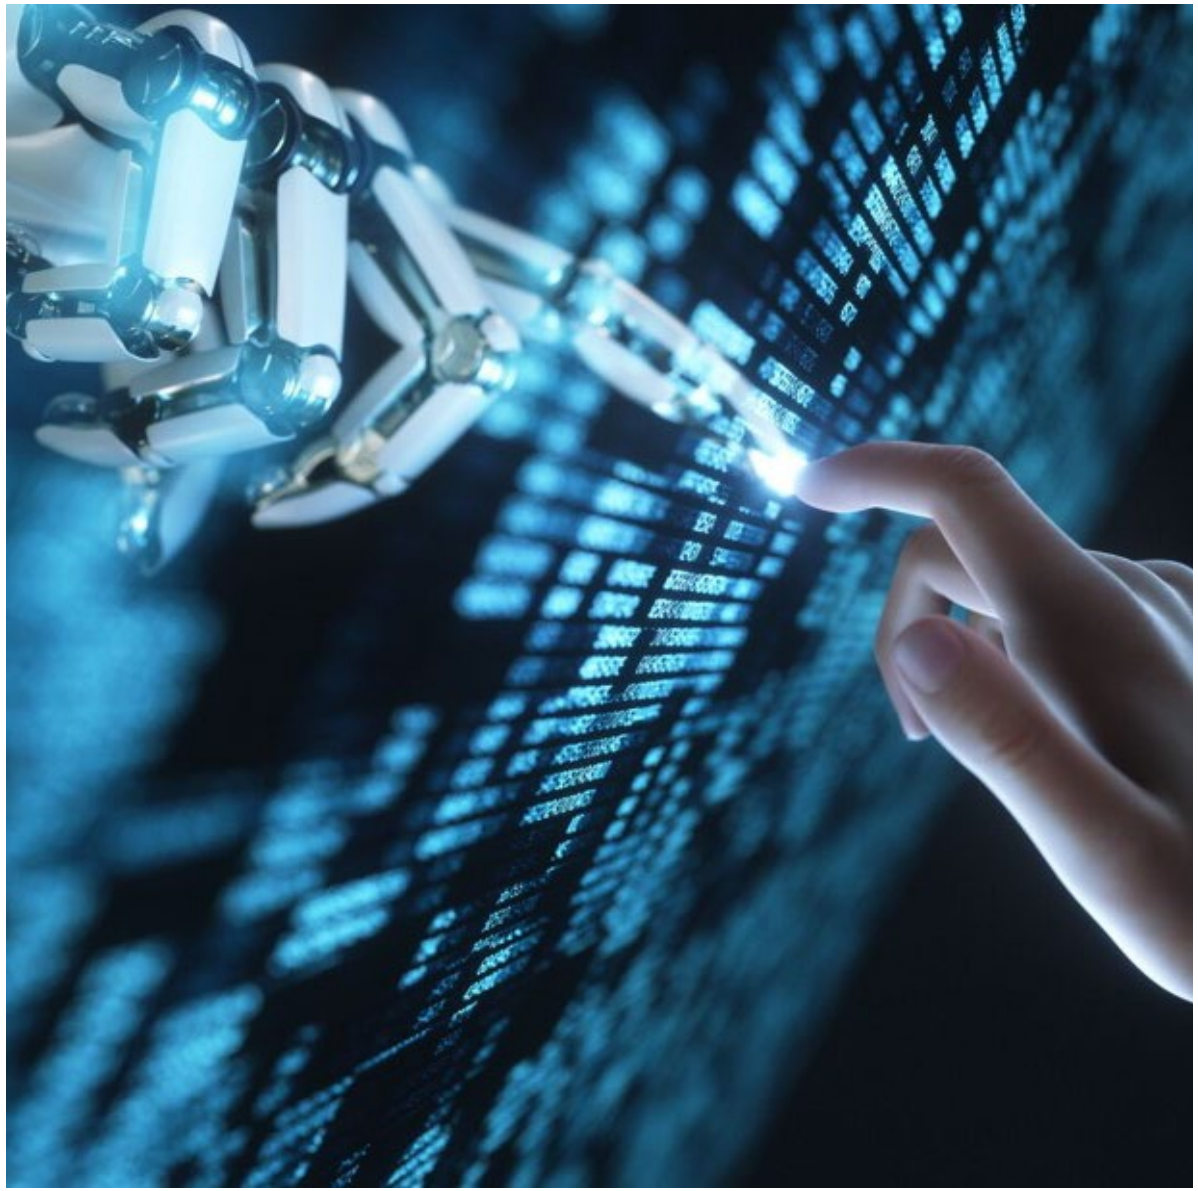
You can’t manage risk you can’t see, and you can’t govern assets you don’t know exist.

First, catalog what exists:

- Agent inventories, hidden agents, embedded AI functions, and shadow AI
- Sanctioned copilots, internal agents, AI APIs, plugins, SaaS AI features, and low-code automations

Then, capture the details:

- Owner, business process, data accessed, connected tools, identity used, and approved actions
- Whether the agent acts as a user, service account, application identity, or delegated workflow
- High-risk combinations: sensitive data + external content + tool access + unattended execution



© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

Implementing AI Responsibly
AI Governance & Risk Management

1 AI Governance

- Develop clear AI policies & procedures
- Establish a governance structure with defined roles
- Identify & manage AI risks
- Promote transparency & fairness
- Implement security controls
- Align AI use with ethical standards

2 NIST AI RMF

- Apply NIST guidelines to manage AI risks
- Define & uphold AI trustworthiness & security
- Support compliance & follow regulations
- Promote responsible AI practices

3 AI IT Audit

- Assess AI model implementation
- Analyze data handling & security
- Review governance effectiveness
- Recommend improvements
- Provide assurance to stakeholders

© 2026 Forvis Mazars, LLP. All rights reserved.

Extending Cyber Resilience to AI
The Agentic Dependency Trap

Consider what vendors and models should be added to a risk register.



81%

of executives say a seven-day outage at their primary AI vendor would have a severe or critical impact, according to IBM IBV.



91%

of executives say they do not have an excellent understanding of their AI dependencies across vendors, models, and infrastructure, per IBM IBV.



71%

of executives say switching their primary AI vendor or model would be difficult, according to IBM IBV.

Source: “The calculus of AI sovereignty,” ibm.com, June 17, 2026.

© 2026 Forvis Mazars, LLP. All rights reserved.

Extending Existing Capabilities
AI Incident Response & Resilience

Prepare & Monitor

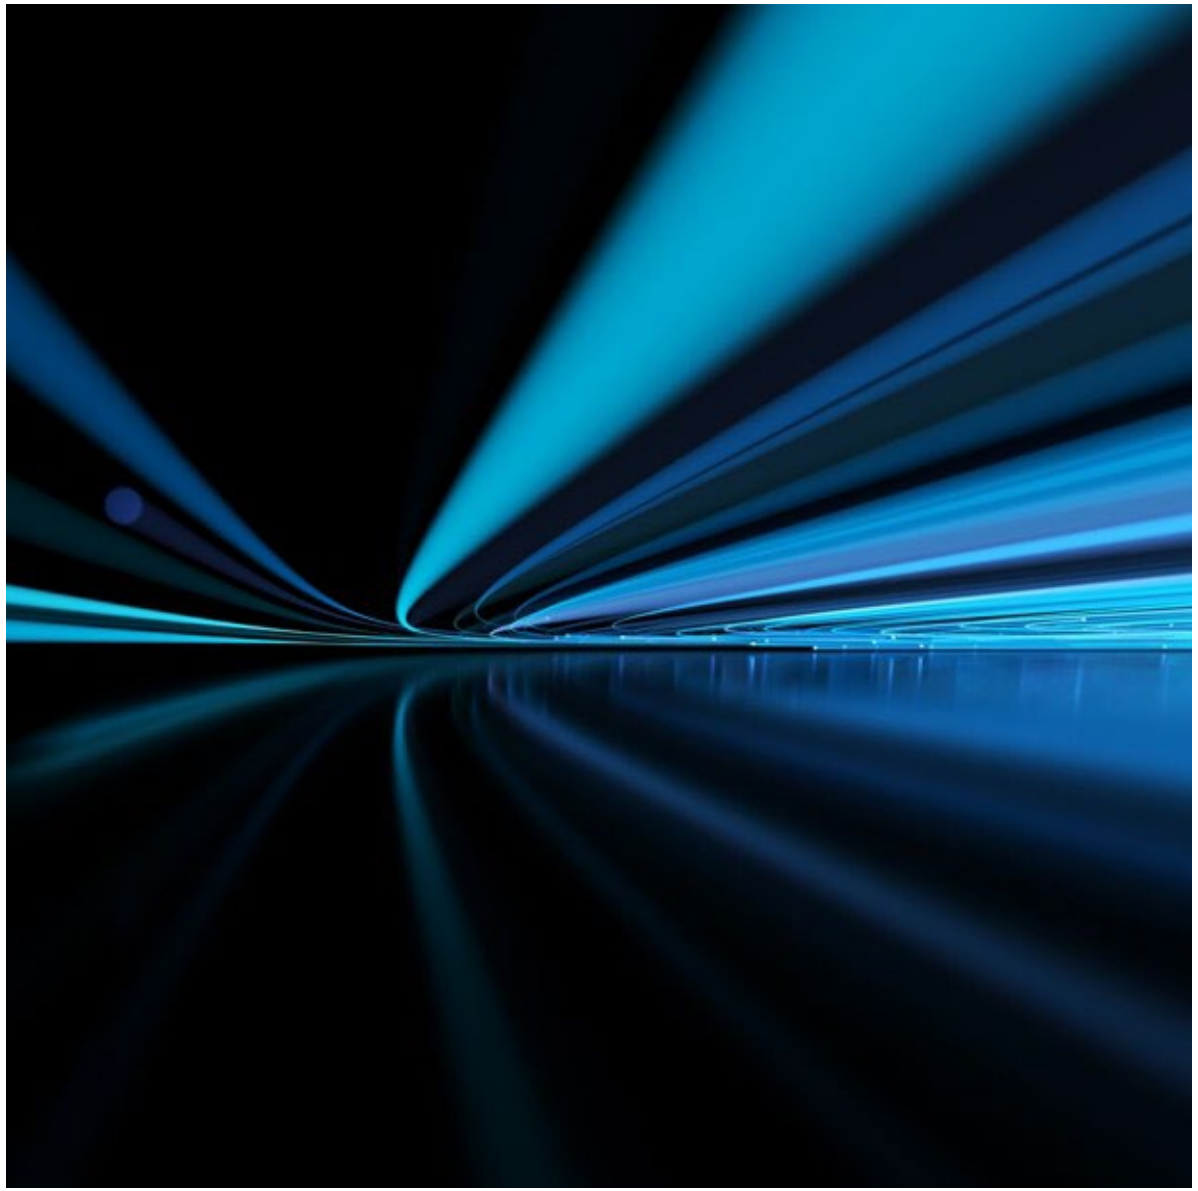
- Define AI-specific incident scenarios and escalation criteria
- Monitor agent activity, model behavior, access, and dependencies
- Retain logs that support investigation and traceability

Contain & Respond

- Disable affected agents, tools, or workflows
- Restrict access and isolate compromised data or integrations
- Activate manual processes, alternate models, or vendor fallbacks

Recover & Adapt

- Validate data, configurations, and outputs before restoration
- Restore services through tested recovery procedures
- Update controls, monitoring, and response playbooks



© 2026 Forvis Mazars, LLP. All rights reserved.

Market Signals, SaaS Trends, & Sovereignty

Vendor concentration, data residency, and model dependency are becoming board-level questions, not procurement details.



Signals From the SaaS Market
Concentration & Dependency

AI Is Concentrating the Market

- AI drove 81.7% of SaaS M&A value and 37.7% of deal count in H1 2026
- M&A hit a decade high of $439.7B, 88.7% from strategic acquirers
- The five largest deals came from SpaceX, Alphabet, NVIDIA, and IBM

Valuations Are Resetting to Fundamentals

- Median SaaS PE EV/EBITDA fell to 11.7x from 20.4x, near non-SaaS at 11.0x
- Software fell below 15% of U.S. syndicated LBO loans, from ~35% in 2024
- Defensible profiles cluster in AI, AI-enabling, and cybersecurity

Why It Belongs on the Risk Register

- Fewer acquirers of AI capability means rising vendor concentration
- Capital rewards AI differentiation, not AI experimentation
- Dependency, portability, and residency now sit alongside cost



Source: PitchBook, “State of SaaS,” H1 2026; global data as of June 30, 2026.

© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

Key Takeaways

Treat each AI agent or agentic AI workflow like a new hire: provide a job description, a manager, and a performance review.

Autonomous risk is operational risk

Agentic AI introduces new exposure as autonomous systems access data, interact with systems, and initiate actions.

Govern the action

Controls must reach the actions agents take, not just the content they produce, making them owned, least-privileged, monitored, and able to be stopped.

Know what you depend on

Map the agents, vendors, and models you rely on, then build inventories, apply risk-based guardrails, and extend incident response to AI.

© 2026 Forvis Mazars, LLP. All rights reserved.



IT Risk & Compliance

Questions?

© 2026 Forvis Mazars, LLP. All rights reserved.



Don’t Miss the Next Session of the

2026 Cybersecurity Virtual Symposium

AI Governance: Turning Accountability Into Action

Tomorrow at 10:00 a.m. - 11:00 a.m. ET

© 2026 Forvis Mazars, LLP. All rights reserved.

forvis mazars

Contact

Forvis Mazars

Sean Andrews
Manager, IT Risk & Compliance
P: 703.970.0443
sean.andrews@us.forvismazars.com

Ricardo Martinez
Partner, Technology, Software & Devices
P: 669.200.2052
ricardo.martinez@us.forvismazars.com

The information set forth in this presentation contains the analysis and conclusions of the author(s) based upon his/her/their research and analysis of industry information and legal authorities. Such analysis and conclusions should not be deemed opinions or conclusions by Forvis Mazars or the author(s) as to any individual situation as situations are fact-specific. The reader should perform their own analysis and form their own conclusions regarding any specific situation. Further, the author(s)' conclusions may be revised without notice with or without changes in industry information and legal authorities.

© 2026 Forvis Mazars, LLP. All rights reserved.