From Identity Security to AI Governance

Why access control is becoming a board-level AI risk

Shawn McCabe | Director of Sales Engineering, US Public Sector, Delinea

Brian McHenry | Vice President, Sales Engineering, Delinea



What we will cover

With short pauses for your questions along the way

01 Why identity is now part of AI governance

02 How AI changes privileged access risk

03 What boards need to understand

04 Where to start

AI agents are identities too

Why identity security is becoming the core of AI governance

AI has moved from answering to acting

Each step gives AI more access, and puts a human further from the decision

Assistants: answer

- **Access:** What the user pastes in
- **Main risk:** Sensitive data leaking out

Copilots: assist

- **Access:** Everything the user can reach
- **Main risk:** Oversharing at machine scale

Agents: act

- **Access:** Their own credentials, often privileged, across systems
- **Main risk:** Actions nobody explicitly approved

The AI enforcement gap

Organizations have the policies. What they lack is enforcement at the moment an agent acts.

- **Policy is not the problem:** almost every organization has a formal policy on what AI can access
- **Enforcement is:** only about half check AI access against that policy in real time
- **AI goes out of bounds:** most have seen AI reach sensitive data beyond its intended scope
- **Detection lags:** few spot a scope violation as it happens; many take a day or longer
- **Accountability is thin:** most cannot reliably trace AI access back to a named human approver
- **Pressure runs the wrong way:** security is being pushed to loosen identity controls for AI

Source: Delinea, 2026 Identity Security Report: The AI Enforcement Gap (Sept. 2026) and Uncovering the Hidden Risks of the AI Race (Mar. 2026).

Governing the model vs. governing what it can touch

Most AI governance starts with the model. Most AI risk shows up in the access.

What AI governance usually covers	What attackers and auditors will ask
Model selection and approval	Which identities can the AI use?
Bias, fairness and explainability	What can it read, write, approve or move?
Data use and privacy policy	Whose credentials, and for how long?
Acceptable-use guidelines	Who is accountable when it acts?

Why AI identities break traditional IAM

The principles still apply: least privilege, ownership, review. The processes do not.

Scale

Agents are created in minutes by developers and business teams. Joiner-mover-leaver processes never see them.

Speed

An agent can act thousands of times before a person notices. Exposure windows shrink from days to minutes.

Delegation

Agents act on behalf of people, often with copied permissions. Did the person act, or the agent?

Q/A

Questions so far?

Identity as part of AI governance

Privilege, with an agent in the loop

How access risk changes when AI connects to enterprise workflows

Scenario: the helpful agent

A bank deploys an AI agent to speed up payment exception handling. It works brilliantly.

It reads customer and transaction data across three systems

It updates cases and drafts approvals in the workflow tool

It runs on a service account created during the pilot, 18 months ago

Its permissions were cloned from a senior analyst "to avoid blockers"

Nobody has reviewed that access since go-live

Nothing here is unusual. Every decision made sense at the time.

Four ways the helpful agent goes wrong

None of these need a sophisticated attacker. All of them use access the agent already has.

Manipulation

A case note contains hidden instructions. The agent follows them, using its real permissions.

Credential theft

A long-lived key in a config file is stolen. The attacker inherits access to all three systems.

Excessive agency

Cloned permissions mean it can approve, not just draft. One bad decision executes.

Toxic combinations

Read customer data + update cases + approve payments. Segregation of duties quietly disappears.

Who finds out, and what good would have looked like

The activity looks legitimate, so security is rarely first to know

Who usually finds out	What good would have looked like
Internal audit, sampling access	A named business owner from day one
Operations, chasing a reconciliation gap	Its own identity, scoped to read and draft
A customer, through a complaint	Short-lived credentials, not standing keys
A regulator or third party	Human approval above a payment threshold
	Included in access reviews and monitoring

Authentication is not authorization

Checking an agent at the door says nothing about what it does once inside

Access at the door	Authorization at the moment of action
Confirms who the agent is, once	Every command, query or tool call evaluated
Agent keeps whatever access it inherited	Out-of-policy actions stopped before they run
Long-lived credentials the agent holds	Just-in-time credentials the agent never sees
Nothing checks what it does next	Every action recorded and attributable

Runtime authorization controls what an agent does.

Session recording proves it.

The Delinea approach

Identity as the control plane for AI

Extend proven privileged access controls to every AI identity

1. Discover every agent, what it can reach, and who owns it
2. Assess risk continuously; flag inherited and excessive privilege
3. Authorize every action at runtime, not just the login
4. Zero standing privilege: just-in-time credentials the agent never sees
5. Prove it: record every privileged action, traceable to a named human

Q/A

Questions so far?

AI and privileged access risk

Making it a board conversation

Accountability, regulation, and access as business risk

Three questions every board should be able to ask

And expect answers backed by evidence, not reassurance

01

What can our AI systems access, and on whose behalf?

02

Who is accountable when an AI system takes an action?

03

How quickly can we see, limit or revoke that access?

Regulators are already asking these questions

Examples from financial services: existing rules that already reach AI access

Framework	What it expects	What it means for AI access
DORA (EU)	Management body accountable for ICT risk; access limited to what is needed	AI agents are ICT assets; their access needs the same controls
NYDFS Part 500	Limit and review privileged access at least annually; MFA; informed board	AI service accounts fall under privileged access rules
EU AI Act	High-risk uses, such as credit scoring, need human oversight and cybersecurity	Access controls are how human oversight is enforced
SR 11-7 / PRA SS1/23	Model inventory, ownership and validation	Extend the inventory to the identities models use
UK SM&CR	Named senior managers accountable for their areas	Someone specific owns AI-driven outcomes

AI risk lives in the handoffs

Four teams each own part of the problem. Identity is the thread that connects them.

Cyber & IAM

Owns: identities, privilege, credentials

Needs: visibility into every AI agent and its access

Data

Owns: classification and data access

Needs: to know which AI systems touch sensitive data

AI & engineering

Owns: models, agents and integrations

Needs: guardrails that don't block delivery

Risk & compliance

Owns: frameworks, attestation, the board report

Needs: evidence that controls actually work

Access control is a business risk. Say it that way.

Same facts, different language. Which version gets a board decision?

Standing privilege

- **Technical:** "Service accounts hold standing admin rights."
- **Business:** "Accounts nobody owns could move money or expose customer data."

AI agent credentials

- **Technical:** "Agents use long-lived API keys."
- **Business:** "If one key leaks, an attacker can do anything the AI can, at machine speed."

Access reviews

- **Technical:** "Quarterly access reviews are behind."
- **Business:** "We can't prove to regulators who had access to what, or why."

Where to start: the next 90 days

No new budget line is needed to begin. It starts with asking the right questions.

Days 1-30: Discover

Inventory AI agents and the identities they use

Assign a named owner to each

Rank them by proximity to money and customer data

Days 31-60: Contain

Remove cloned human permissions

Replace long-lived keys with short-lived credentials

Add human approval for high-impact actions

Days 61-90: Govern

Add AI identities to access reviews

Include an AI-agent scenario in incident exercises

Report AI access risk to the board

THE ONE THING TO TAKE BACK

Every AI agent is a privileged identity

Give it a named owner. Limit it to the task.

Make its access visible to the board.

Q/A

Your questions

Thank you for listening

Thank you.

Delinea

Securing identities at every interaction

Further reading: Delinea 2026 Identity Security Report, "The AI Enforcement Gap"
delinea.com/resources/ai-policy-enforcement-report

Continue the conversation:

Shawn McCabe
Director of Sales Engineering, US Public Sector
shawn.mccabe@delinea.com

Brian McHenry
Vice President, Sales Engineering
brian.mchenry@delinea.com