




AI attacks are now a board risk

A CISO brief on protecting revenue,
trust, and resiliency



EXECUTIVE SUMMARY

Why the attack path runs through identity—and why resilience requires controls that act at runtime.

The economics of cyberattacks have changed.

AI-powered attackers can discover risks faster, adapt continuously, and pursue multiple attack paths in parallel. Once inside an environment, they can use legitimate credentials, identities and privileges to move laterally and escalate access, often without exploiting a new vulnerability at every step.

That changes the security question. The critical issue is no longer simply whether an attacker can get in. As AI makes footholds faster and more likely, the question is now how far an attacker can move once they do and whether that foothold can become a business-impacting incident.

For most organizations, that question is especially important. Enterprises operate complex environments spanning Active Directory (AD), legacy systems, and homegrown applications. Much of this infrastructure cannot be modernized overnight—but **AI attacks will not wait for modernization to finish.**

Assume the foothold will happen: Contain movement before a compromise becomes a business-impacting incident.

A resilient architecture needs a control that can evaluate access in real time and, when risk warrants it, challenge or block the request before a foothold becomes lateral movement, privilege escalation, or a broader business-impacting incident.

For the board, four implications follow:

- **Identity is the attack path.** AI attacks turn legitimate identities, credentials and privileges into movement through the environment.
- **Preventing a foothold cannot be treated as the end of the story.** Post-compromise containment must become the primary resilience objective.
- **Modernization alone cannot be the near-term answer.** Replacing legacy infrastructure takes years and can require significant investment and business disruption. **Resilience cannot wait for the modernization program to finish.**
- **Controls must operate at attack speed.** The ability to intervene at the point of access can stop lateral movement and privilege escalation and reduce the blast radius while broader remediation continues.



In a Fortune 50
red team exercise using
Anthropic's Mythos,

an autonomous AI-powered attacker reached domain compromise in roughly two hours by chaining identities and access paths.

Silverfort stopped the AI attacks inline, consistently and where all other defenses failed - demonstrating that runtime Identity Security is a critical control plane for resilience against these kind of attacks.

Why every organization is exposed

Organizations have spent years building layered security around increasingly complex environments. **That complexity is now a resilience challenge.**

Active Directory, legacy systems and authentication protocols, and homegrown applications create a dense web of relationships between users, machines, applications and systems.

AI can accelerate the process of discovering and exploiting those trust paths while automating reconnaissance, credential use, privilege escalation, and lateral movement across multiple systems.

That creates a new operational reality: Attacks can be completed at machine speed, before there is time to detect, investigate, contain, and respond.

What the board is actually protecting

For any organization, the consequence of an AI attack is not limited to compromised systems or stolen data. The board is ultimately protecting revenue, customer trust, operational continuity, and the institution's ability to operate through disruption.



Revenue and financial performance

Service disruption, downtime, and customer attrition can create material financial losses.



Operational resilience

Core business systems and critical applications cannot simply be taken offline while teams investigate and recover.



Reputation and customer trust

A serious compromise can undermine confidence, particularly when critical services are disrupted or a breach becomes public.



Stakeholders & regulatory confidence

Demonstrable resilience matters to regulators, customers, investors, and markets, particularly the ability to contain attacks and continue operating through disruption.

The questions the board should be asking:

If an attacker gained a foothold today, how far could they move through our identities?
Do we have the control points and established processes needed to stop an attack?

Traditional security metrics, such as vulnerabilities, misconfigurations, alerts and control coverage, remain important, but the board needs to know something more fundamental: **Can we contain an AI attack before it becomes a business-impacting incident?**



Foothold



Credential



Authentication



Privilege escalation



Lateral movement



Critical systems

WHAT IT MEANS FOR
SECURITY LEADERS:

**What stops
the next step?**

Every transition creates an access decision. **That is the control point.**

Ask yourself: *If a compromised identity requests access to the next system, can we evaluate and control that request in real time, before the attacker can use it to reach the next identity, privilege or system?*

If the answer is no, there may be a gap between detecting an attack and actually containing it.

Three assumptions security leaders should reconsider

1. “If we patch faster, we stay ahead.”

Patching is essential, but AI attacks exploit existing weaknesses, legitimate credentials, and excessive privileges. The attacker does not necessarily need a new vulnerability. Patch faster, but don't assume patching alone will stop the attack path.

2. “We already have IGA and PAM.”

IGA and PAM remain essential for governance, operational control, and compliance, but they are too static and limited. Resilience means continuously evaluating and controlling access as an attack unfolds, including when legitimate identities are compromised.

3. “Modernization will solve the problem.”

Modernization is necessary, but it takes years. Meanwhile, frontier AI capabilities are advancing rapidly, compressing attack timelines and increasing the urgency to contain attacks in today's environments. Modernize for the future. Build resilience for the attacks happening now.

The new control point: identity

Once an attacker has a foothold, many of the next steps in an attack require authentication or authorization. That creates an opportunity to intervene.

Before access is granted, evaluate the request. Block the risky access while keeping the rest of the business running.

The objective is not to replace previous investments in other tools (EDR, IGA, PAM, vulnerability management). It's to add a control point that those technologies do not inherently provide and that has been proven effective in building resilience against AI attacks: **inline, runtime enforcement at the moment of access.**

Inline, runtime identity enforcement can help organizations:

- **Challenge** risky authentication
 - **Block** lateral movement
 - **Prevent** privilege escalation
 - **Contain** compromised identities
 - **Restrict** access to critical systems
 - **Reduce** the blast radius of a successful foothold
-

Regulatory pressure is moving in the same direction

The regulatory framework is still evolving. The supervisory expectation is not.

AI-powered cyber risk is moving from a technical concern to a financial and operational resilience issue.

In June 2026, the cybersecurity agencies of the Five Eyes nations jointly warned that frontier AI models could transform offensive cyber capabilities in months, not years, putting new urgency on how organizations defend against AI-powered attacks. They also called on business leaders to treat the implications as a core governance responsibility.

That same month, the European Systemic Risk Board warned that frontier AI could dramatically increase the speed, scale, and sophistication of cyberattacks, shrinking the time available for defenders to respond.

In July 2026, the European Commission launched an Action Plan on Cybersecurity and AI, highlighting how advanced AI can automate attacks and increase the scale and speed of cyber incidents.



**The message
for the board:**

Don't wait for the rulebook to catch up. Institutions increasingly need to demonstrate that they understand their exposure, can respond at machine speed, and can contain an attack before it becomes a business continuity event.

What slow action already costs

A serious breach creates costs far beyond technical remediation.



Business operational cost

Lost production, missed commitments, business downtime, and prolonged business interruption



Direct financial cost

Revenue loss, customer attrition, emergency spending, and the cost of restoring operations.



Escalating cost of exposure

The longer compromised identities and systems go undetected, the greater the eventual scope—and cost—of damage and recovery.



Legal and regulatory cost

Fines, litigation, breach disclosure, regulatory response, and mandated remediation.



Response and recovery cost

Incident response, containment, forensic investigation, and sustained strain on IT and security teams.



Unplanned modernization cost

A breach can force organizations to replace legacy infrastructure, and accelerate capital spending on an attacker's timeline.

The cost of waiting means higher risk and fewer options when an attack begins.

\$4.99M

Global average cost of a data breach, up 12% year-over-year

~2 hours

For an AI attacker to reach domain compromise in a Fortune 50 red team exercise

+\$1M

Additional cost for AI-driven attacks

PROOF, NOT THEORY:

A Fortune 50 Mythos red team exercise

A Fortune 50 organization tested what an autonomous AI attacker could accomplish inside its environment. **The results:**

- **~2 hours** to reach domain compromise
- **7 identities** chained through the attack path
- **Silverfort** stopped the attack path inline

The attack was able to move through legitimate identities and access relationships at machine speed. Silverfort's inline identity controls interrupted the attack path. **The team ultimately had to disable the controls to allow the exercise to continue.**

“Silverfort is phenomenal. It blocked Mythos' attempts to spread in the network. At some point, we had to disable Silverfort's defenses to allow further testing.”

— Security operations leader,
Fortune 50 company

Sources: [Silverfort – the Mythos field report](#); [IBM 2026 Cost of a Data Breach Report](#).

What the board should decide this quarter

01. Establish recurring readiness assessment

Ask for a board-visible, recurring readiness assessment —not an annual audit—that quantifies exposure to AI attack chains across core business systems, critical infrastructure, and vendor systems.

02. Make AI attack resilience a strategic priority

Treat AI attack defense as a growth, trust, and resilience investment, not a compliance line item.

03. Establish runtime identity enforcement as a key resilience control

Apply runtime, inline identity enforcement as a control plane for resilience, bringing human, service, non-human and AI-agent identities under control across all environments, without waiting for multi-year modernization projects.



Silverfort: Stop attack paths before they impact the business

Silverfort helps organizations put a **runtime identity control point** in place quickly across existing environments.

Proven in production across 1,000+ organizations worldwide, Silverfort enables organizations to:

- **Prevent footholds from becoming security incidents** by blocking lateral movement and controlling identity and authentication flows at runtime to contain attacks.
- **Protect critical business systems rapidly** across legacy and modern environments, without requiring slow and expensive modernization or refactoring.
- **Turn exploitable identity weaknesses into controlled risk** rather than waiting for every underlying issue to be remediated.

THE BOTTOM LINE

Build resilience now. Modernize over time.

AI has changed the economics of cyberattack. Attackers can use the identities, credentials and access already inside the environment and move faster than traditional security processes can respond.

For organizations operating complex, legacy environments, the board needs a new answer to an old question:

How far could an attacker move and what would stop them?

Silverfort helps organizations answer that question and contain attack paths before they become business-impacting incidents.

About Silverfort

Silverfort secures every dimension of identity. We deliver end-to-end Identity Security that is easy to deploy and won't disrupt business operations, resulting in better security outcomes with less work. Discover every identity, analyze exposures, and enforce protection inline to stop lateral movement, ransomware, and other identity threats.