CISO FIELD GUIDE

7-step Playbook to Master the **AI Exploitability Gap**

Practical moves to turn AI risk awareness into action - drawn from the 2026 Global AI & Cyber Maturity Report.

Based on NIST CSF 2.0 & NIST AI RMF 1.0

21	16	2,400+
Countries	Industries	Assessments

> *"The security playbook took decades to build. AI can break it in minutes."*

2026 average AI risk maturity: 2.35 / 5.0

Table of Contents

Awareness is up. Action isn’t.

Organizations have never been more aware of AI risk but most still fail to act on it. Across 2,400+ assessments in 21 countries and 16 industries, the 2026 Global AI & Cyber Maturity Report found the same fault line in both NIST frameworks: governance scores highest, while the functions that turn awareness into action — Protect in cybersecurity, Manage in AI — score lowest.

88% of organizations now use AI in at least one business function, yet their average AI risk maturity is stuck at 2.35 out of 5.0 — the “reactive” level. Policies alone don't reduce risk. Without enforcement, they create a false sense of security, leaving organizations exposed while believing they're protected.

The seven steps that follow, drawn directly from Cye’s recommendations, move a security program from knowing the risk to acting on it — built from the foundation up.

AWARENESS VS ACTION ACROSS BOTH FRAMEWORKS

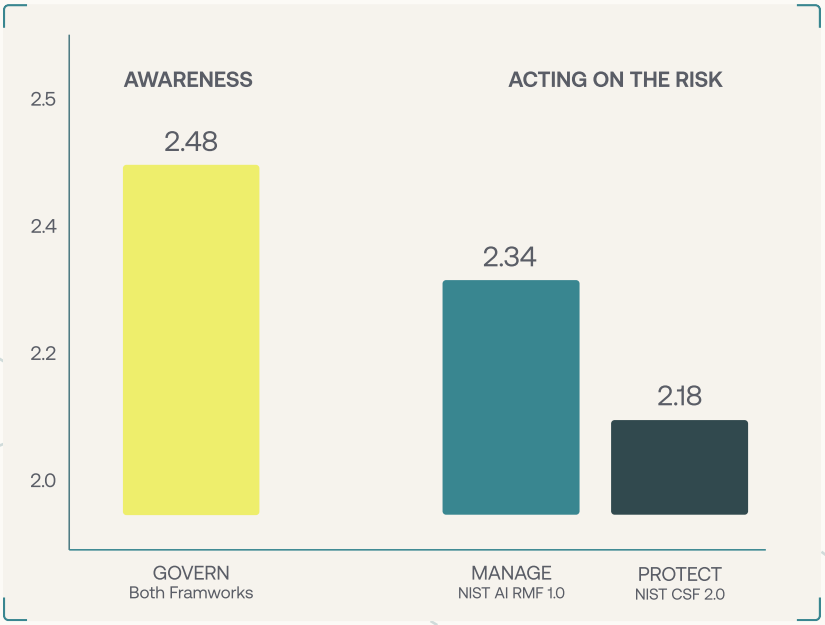

Governance leads in both frameworks; the functions that turn awareness into action fall furthest behind.

Why the Gap Matters

The risk is already inside the enterprise

AI risk isn't a future problem – it's already embedded across sanctioned deployments, employee experimentation, copilots, autonomous agents, and third-party services. Like Shadow IT a decade ago, employees are adopting AI faster than organizations can discover or govern it, and these tools can reach sensitive data, source code, and customer information on their own.

Exposure varies sharply by sector, and the pattern is telling: your bank is safer than your plane or your hospital — not because banks face fewer threats, but because regulation forced them to build the visibility and governance most sectors still lack.

SHADOW AI EXPOSURE BY INDUSTRY

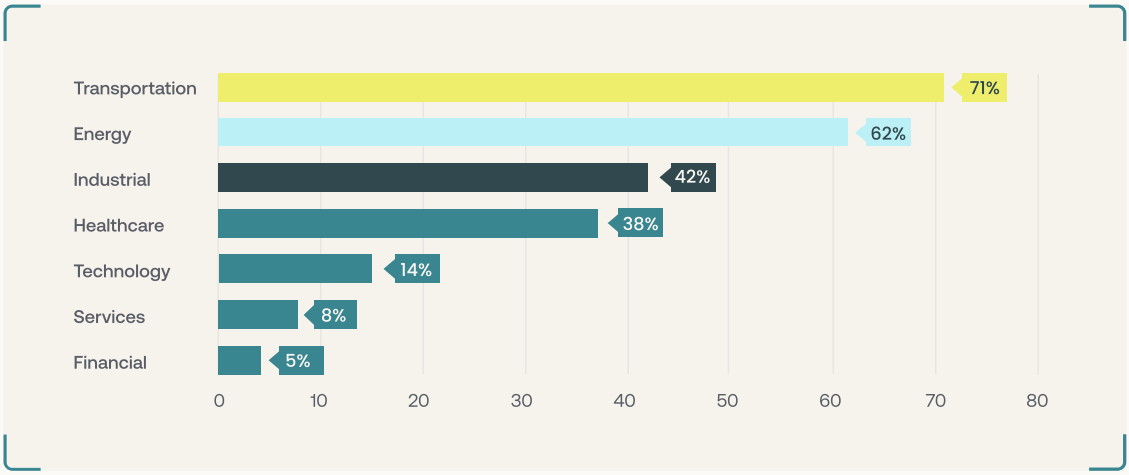

Share of organizations using AI beyond formal governance — below-median AI governance with active AI-related findings.

AI MATURITY BY FUNCTION

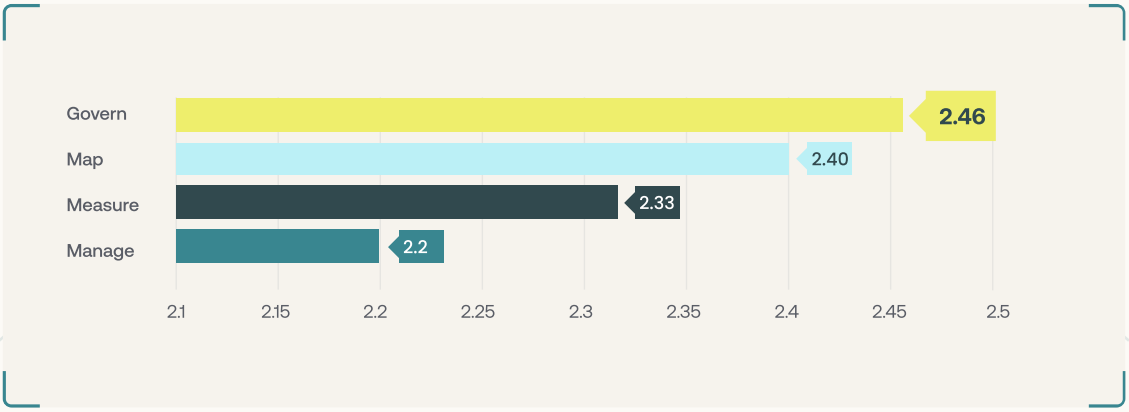

Organizations govern AI risk far better than they manage it — the function that turns awareness into enforcement, response, and recovery.

The Playbook

7 steps to master the gap

Knowing the gap isn't closing itdoesn't reduce your exposure. These steps move organizations from awareness to action, starting with the fundamentals and building upward - each one mapped to the NIST function it strengthens.

STEP 01

Map your AI before your next board meeting

You can't govern what you can't see — and most AI has entered the enterprise without anyone cataloging it: sanctioned tools, embedded copilots, autonomous agents, and the shadow AI employees adopted on their own. Ungoverned AI already reaches 71% of organizations in transportation and 62% in energy. A complete inventory is the single move that unlocks every step that follows.

DO THIS
Build one living inventory of every AI tool, model, API, and service in use — including unsanctioned ones — and tag each with its data sensitivity and owner.

STEP 02

Give AI risk a named owner

The "CISO effect" extends to AI: organizations with explicit executive ownership carry far fewer unresolved AI findings. Without a named owner, AI risk falls into the gap between IT, security, data, and the business — and no one holds the authority to act. Accountability is what turns a written policy into an actual decision.

DO THIS
Assign one accountable executive for AI risk, with a clear mandate, a budget line, and the authority to pause or block deployments.

STEP 03

Fix the basics first

Protect is the lowest-scoring CSF function two years running, at just 2.18 out of 5.0 — and the gap is basic hygiene, not sophistication. The most common findings are outdated technologies, exposed admin interfaces, missing security headers, and insufficient cloud monitoring, all known and fixable for years. AI doesn't change the fundamentals; it just exploits them faster.

DO THIS
Redirect spend to patch management, access controls, and asset hardening before funding anything new — the cheapest, most neglected path to maturity.

STEP 04

Extend vendor risk to AI

Every AI application leans on a chain of external models, APIs, plugins, and data services — and insufficient vendor risk management was the single most common governance finding in the dataset. Most third-party assessment programs were never built to evaluate a model. Every external model is a dependency your program wasn't designed for.

DO THIS
Update vendor assessments for AI-specific criteria — training data, model updates, data residency, and failure modes — and treat each model provider as part of your attack surface.

STEP 05

Monitor AI before you scale it

Most organizations still can't detect AI-driven attacks or misuse, and monitoring blind spots — especially in the cloud — are among the most frequent findings. Scaling AI without visibility simply multiplies the exposure you can't see. You can't respond to what you can't detect.

DO THIS
Instrument for model drift, pipeline integrity, prompt injection, and usage logging before deploying more AI — not after.

STEP 06

Drill the gap, don't document it

The widest divide in the AI framework sits between Govern (2.46) and Manage (2.22) — between writing policy and being able to act on it. Policy alone never closes that distance. The teams that rehearse real failure scenarios consistently outperform the teams that only write them down.

DO THIS
Run tabletop exercises on concrete AI scenarios — prompt injection, data poisoning, a customer-facing chatbot gone wrong — and fix what the drill exposes.

STEP 07

Use regulation as a floor, not a finish line

Regulation, not budget, moved the needle this year: the biggest maturity gains all followed enforced deadlines, with Switzerland jumping +16%. Full EU AI Act compliance is due August 2026, carrying penalties up to €35M or 7% of global turnover — yet most industries still score below what the rules assume. Treat compliance as the baseline to build past, not the goal to reach.

DO THIS
Map your program to upcoming deadlines now, then measure maturity continuously rather than once a year — so awareness keeps turning into action.

The Bottom Line

> *"The organizations that master the gap fastest won't just lower their risk. They'll define what it means to be cyber mature in the age of AI."*
>
> **Nimrod Partush** Chief Innovation & AI Scientist, Cye

About Cye

Cye combines an AI-native exposure management platform with world-class cyber expertise to help organizations know the financial impact of their cyber exposure, prioritize risk mitigation and automate remediation. Cye's 500+ customers gain the clarity to make smart defensible decisions that reduce their risk exploitability with speed, and improve their resilience to the hyper dynamic threat landscape.

Visit us at cyesec.com.