2026 Global AI and Cyber Maturity Report

Cye's Annual Cybersecurity and AI Maturity Review, based on NIST CSF 2.0 and AI RMF 1.0

21
Countries

16
Industries

2,400+
Assessments

Contents

Executive Summary

> "AI is inheriting cybersecurity's oldest problem: the gap between policy and action"

Reuven Aronashvili
Founder and CEO of Cye

Artificial intelligence is breaking traditional cybersecurity faster than organizations can adapt. In the post-Mythos era, where AI can autonomously discover and chain vulnerabilities at scale, attackers can weaponize weaknesses faster than teams can patch them, and time-to-exploit has collapsed from weeks to minutes, discovery is no longer the hard part. The challenge now is separating real risk from noise and acting on it.

This year's report measures overall cybersecurity maturity against NIST CSF 2.0 and, for the first time, AI risk maturity against NIST AI RMF 1.0. The two tell the same story. Both are advancing yet held back by the same weakness: organizations are far better at governing risk on paper than acting on it. AI hasn't created a new problem so much as inherited cybersecurity's oldest one, and amplified it. The real risk in 2026 isn't adopting AI. It's failing to manage the exposure it introduces fast enough to stay secure.

Key Insights

Risk thrives in the gap
between awareness and action.

The same weakness runs through both frameworks. Governance scores highest but the functions that turn awareness into action, Protect in cyber and Manage in AI, score lowest. The evidence is clear: knowing the risk isn't the problem. Acting on it is.

AWARENESS VS ACTION ACROSS BOTH FRAMEWORKS

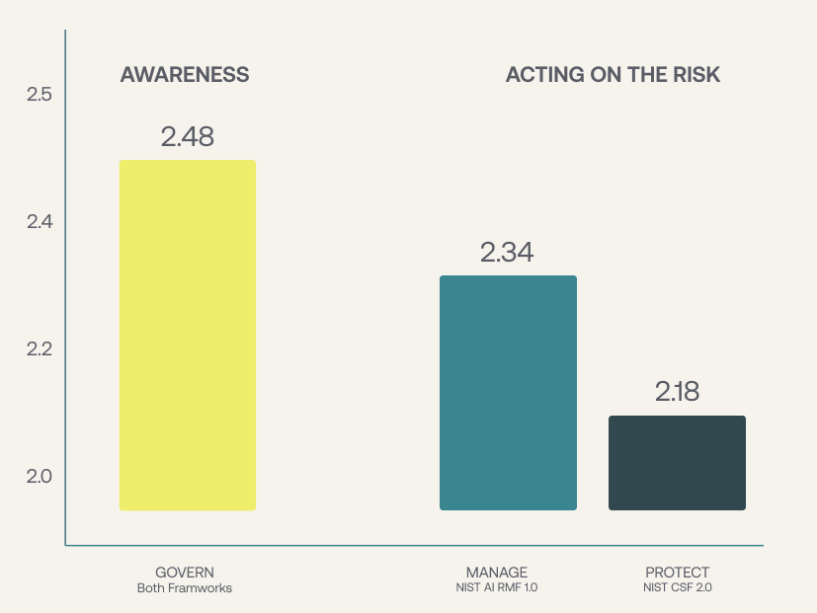

AI adoption has
outpaced AI security.

88% of organizations are using AI (McKinsey: The State of AI in 2025), yet their average AI risk maturity is stuck at the 'reactive' level - scoring just 2.35 out of 5.0.

Shadow AI is out of control
and critical infrastructure is most in danger

Ungoverned AI is spreading beyond visibility, hitting critical infrastructure hardest: 71% exposure in transportation and 62% in energy, versus just 5% in financial services.

71%
exposure in transportation

62%
energy

5%
financial services

SHADOW AI EXPOSURE BY INDUSTRY

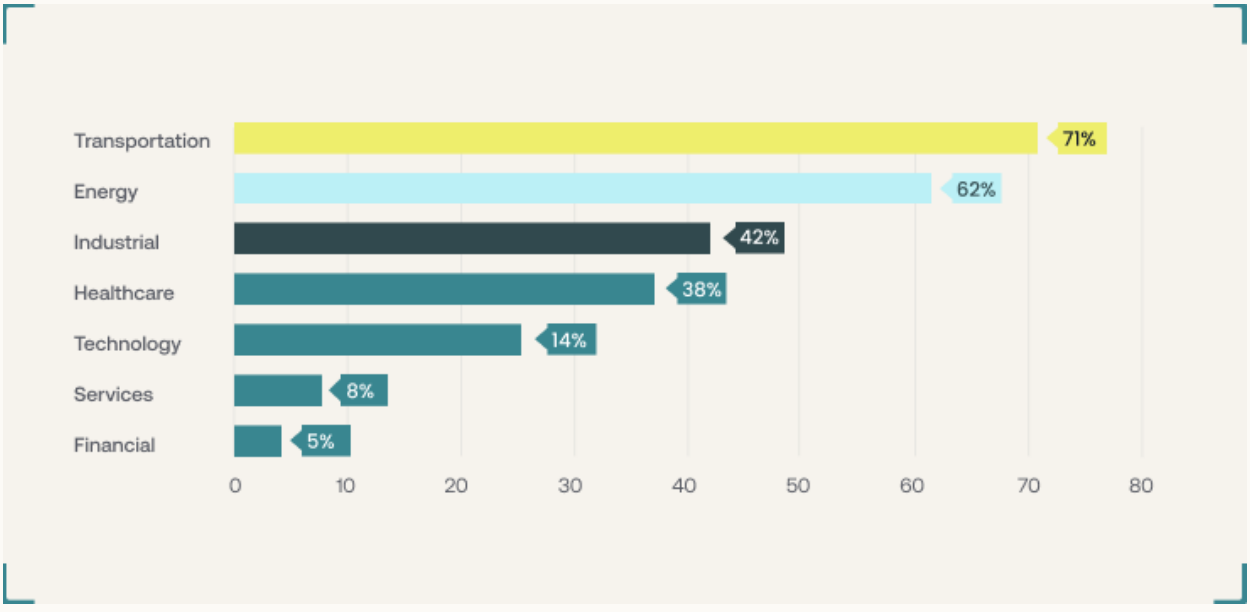

Regulation moves the needle
Budgets don't.

The biggest maturity gains all followed enforced deadlines: Switzerland jumped +16% after two regulations came into force this year, while the heavily regulated financial sector is both the most mature and the least exposed.

WHEN DEADLINES HIT, MATURITY ROSE
MATURITY GAINS BY COUNTRY, 2025 - 2026

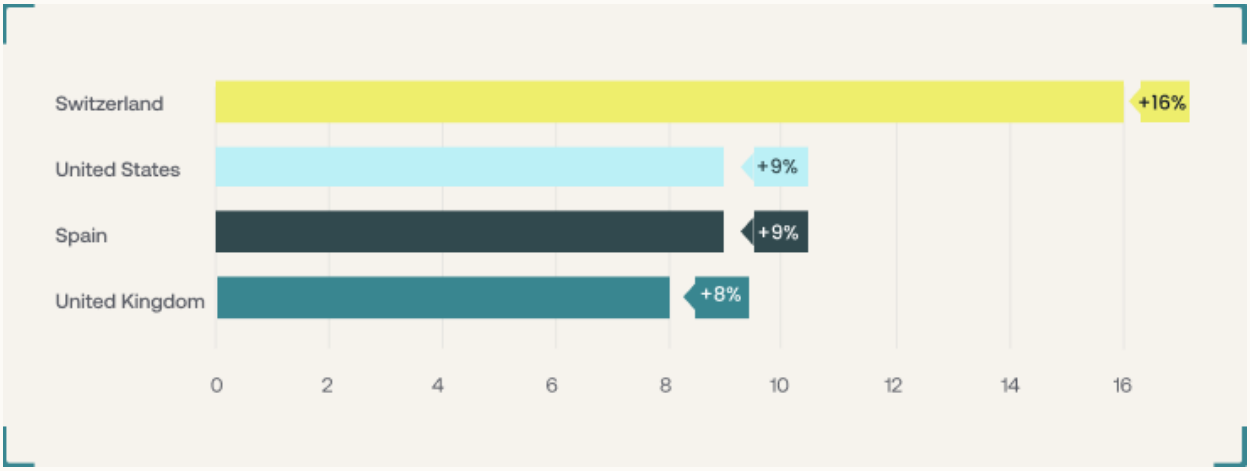

$212B
Global Cybersecurity Spending in 2026

More spending
Lower security.

Global cybersecurity investment hit $212B in 2026, up 15% year over year - yet nearly a third of organizations feel less secure than a year ago. Investment isn't translating into resilience.

The Data Behind The Results

This report draws on more than 2,400 data points spanning 21 countries and 16 industries. We mapped every result to two NIST frameworks, the Cybersecurity Framework (CSF 2.0) and the AI Risk Management Framework (AI RMF 1.0), and measured Shadow AI exposure.

Maturity - NIST CSF 2.0

We scored the six core functions of the NIST CSF 2.0 framework:

GOVERN
Policies, accountability, and risk strategy

IDENTIFY
Asset visibility and risk assessment

PROTECT
Safeguards and access controls

DETECT
Monitoring and anomaly detection

RESPOND
Incident response and communication

RECOVER
Recovery planning and continuity

AI – NIST AI RMF 1.0

We scored the four core functions of the NIST AI RMF 1.0 framework:

GOVERN
AI risk policies, board oversight, accountability

MAP
Identify and catalog AI systems and their risks

MEASURE
Monitor, evaluate and track AI risk continuously

MANAGE
Enforce controls, respond to and recover from AI incidents

Scoring

Every function is scored from 1 to 5, where 1 is the least mature (most vulnerable) and 5 is the most mature.

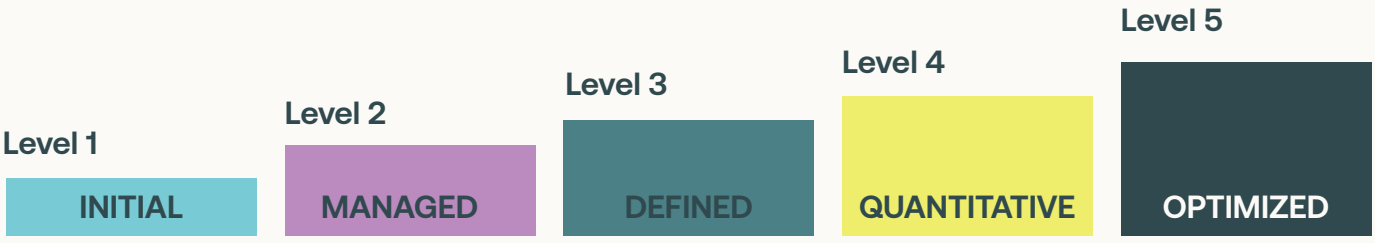

Unpredictable, reactive

Managed at project level, reactive

Managed at org level, proactive

Processes measured controlled

Focus on process improvement

Shadow AI Exposure

Shadow AI is the use of AI tools and systems beyond an organization's visibility or control. We define Shadow AI Exposure as organizations that score below the median on AI governance (median AI GOVERN = 2.66) while also showing AI-related findings: clear evidence that AI is in use but isn't being properly governed.

Part One

The AI Maturity Gap 2026's Defining Risk

Companies Are Racing To Deploy AI - And Attackers To Exploit It

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

Organizations spent years maturing their cybersecurity programs, then wired AI into everything in a fraction of that time without the same discipline. Today 88% use AI in at least one business function (McKinsey- The State of AI in 2025), yet average AI risk maturity is just 2.35 out of 5.0. And almost none of it is fully governed.

That gap is a target. Both sides now hold the same weapon: defenders use AI to detect and respond, while attackers use it to find and exploit weaknesses faster than any human team. Tools like Anthropic's Mythos have already autonomously uncovered thousands of vulnerabilities and chained them into full compromises. Time-to-exploit is now measured in minutes, and the window to close the gap is narrow.

The Governance Illusion

Awareness Is Up, Action Isn't

Organizations know AI introduces risk. They've written the policies and briefed the board. What they haven't done is act on it.

Across the NIST AI RMF, governance scores highest, and management - the ability to enforce, respond, and recover, scores lowest. Organizations are putting in policies that they can’t act on.

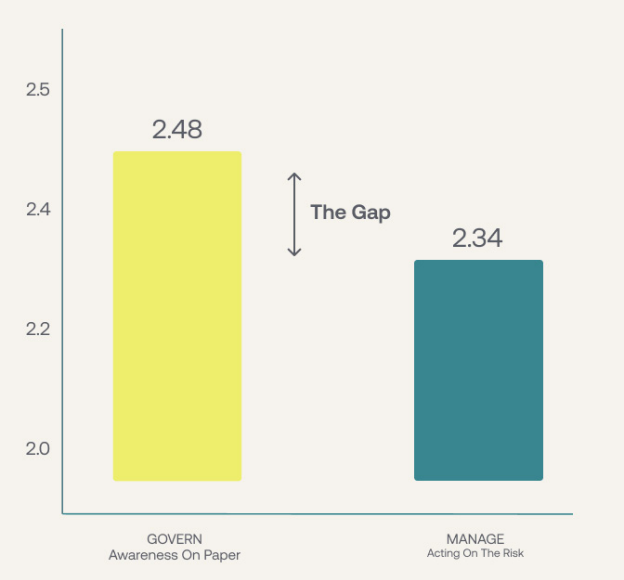

The gap varies sharply by sector

Education and energy fall furthest behind, while even tightly regulated financial services can't fully close the gap.

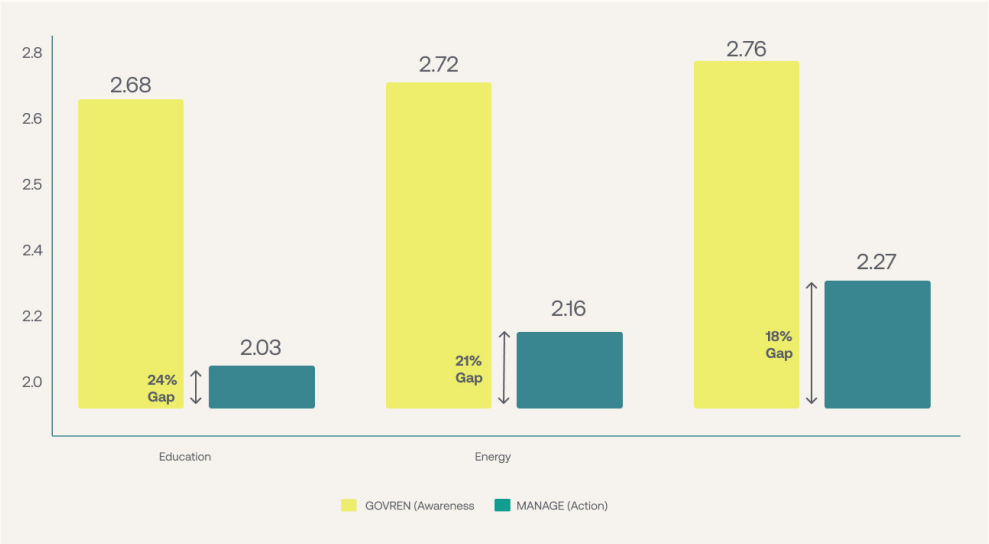

Policy without enforcement is just exposure. And automating AI security decisions without business context only speeds up the wrong fixes.

AI Risk Is Already Inside the Enterprise

This isn't a future problem. AI risk is already live, across sanctioned deployments, employee experimentation, embedded copilots, autonomous agents, and third-party services.

Shadow AI
You Can't Govern What You Can't See

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 wildly by sector, and the pattern is telling: your bank is safer than your plane or 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

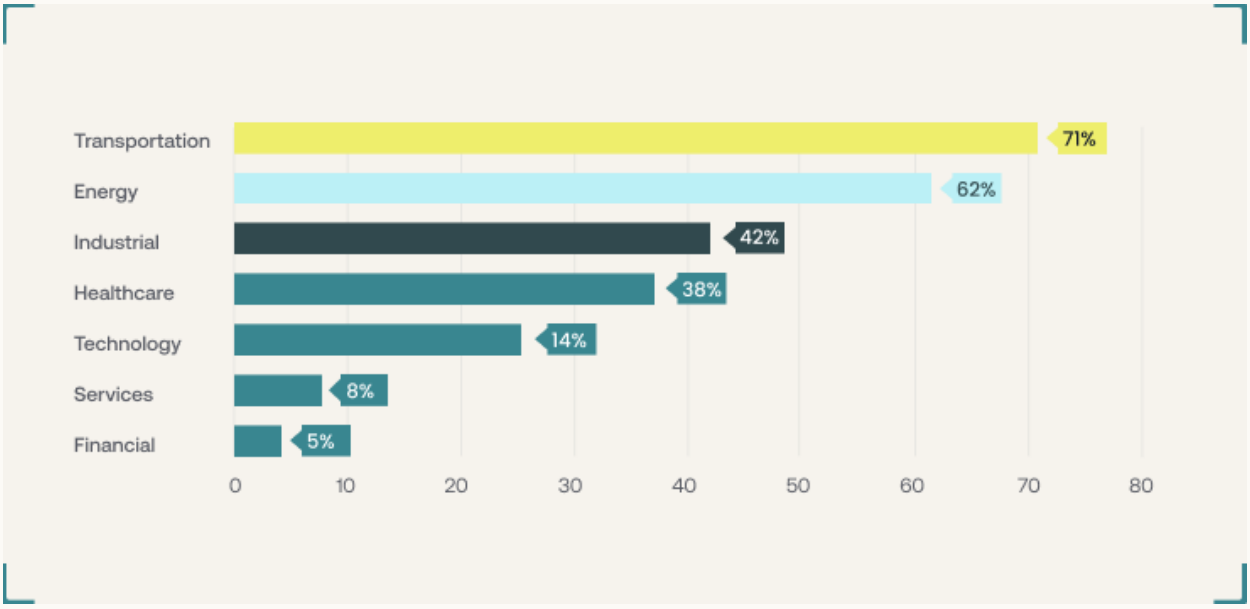

AI Risk Is Already Inside the Enterprise

AI Risk is Active

We identified 134 finding types tied directly to AI systems in live production environments. The largest share are AI infrastructure misconfigurations, alongside identity and access gaps and monitoring blind spots that let misuse go undetected.

SEVERITY OF ACTIVE AI FINDINGS IDENTIFIED

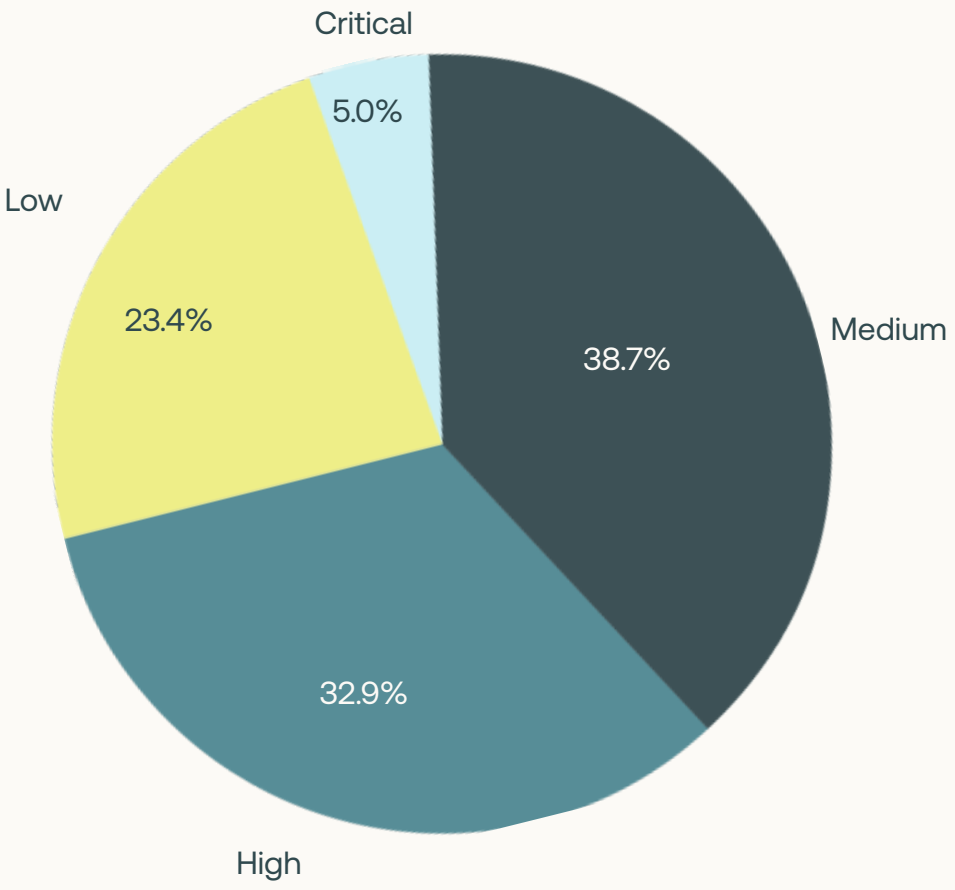

> "The AI attack surface is already here, in the tools you've sanctioned, the ones you haven't, and a supply chain you can't fully see."

AI is the New Supply Chain Risk

Every AI application leans on a chain of external models, APIs, plugins, and data services. The single most common governance finding was insufficient vendor risk management, and most programs were never built to assess it.

Nimrod Partush
Chief Innovation & AI Scientist, Cye

AI Risk Up Close
By Function, Industry, and Country

AI RMF 1.0 – Breakdown by Function

Across the four AI RMF functions, one weakness stands out: organizations govern AI risk far better than they manage it, the function that turns awareness into enforcement, response, and recovery.

AI MATURITY BY FUNCTION

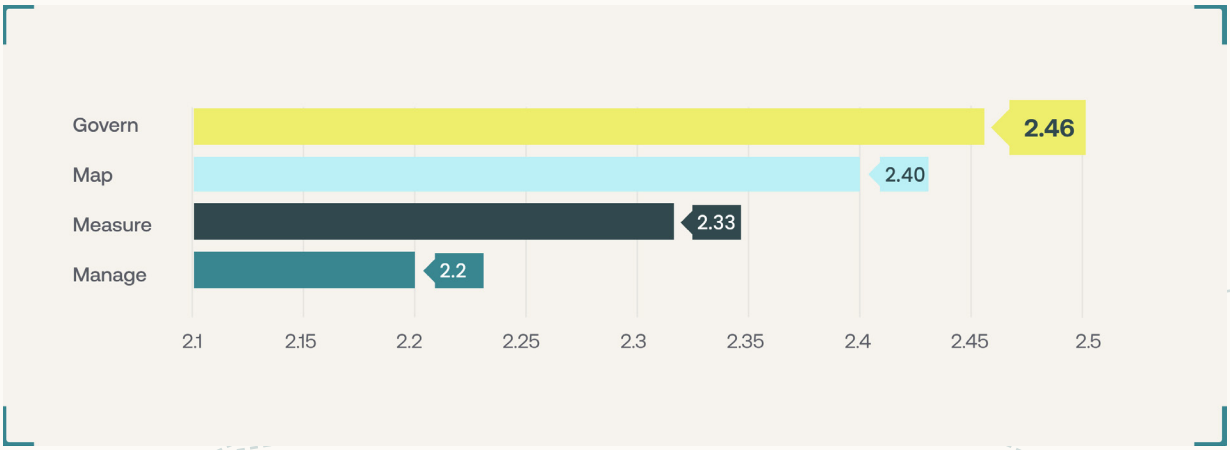

No Industry is AI-Ready

No industry has closed the AI maturity gap.

Financial and Technology lead the field, but even the front-runners fall short of mature, and consumer and transportation trail badly. The entire field is stuck low in the "Managed" band.

AI MATURITY BY INDUSTRY

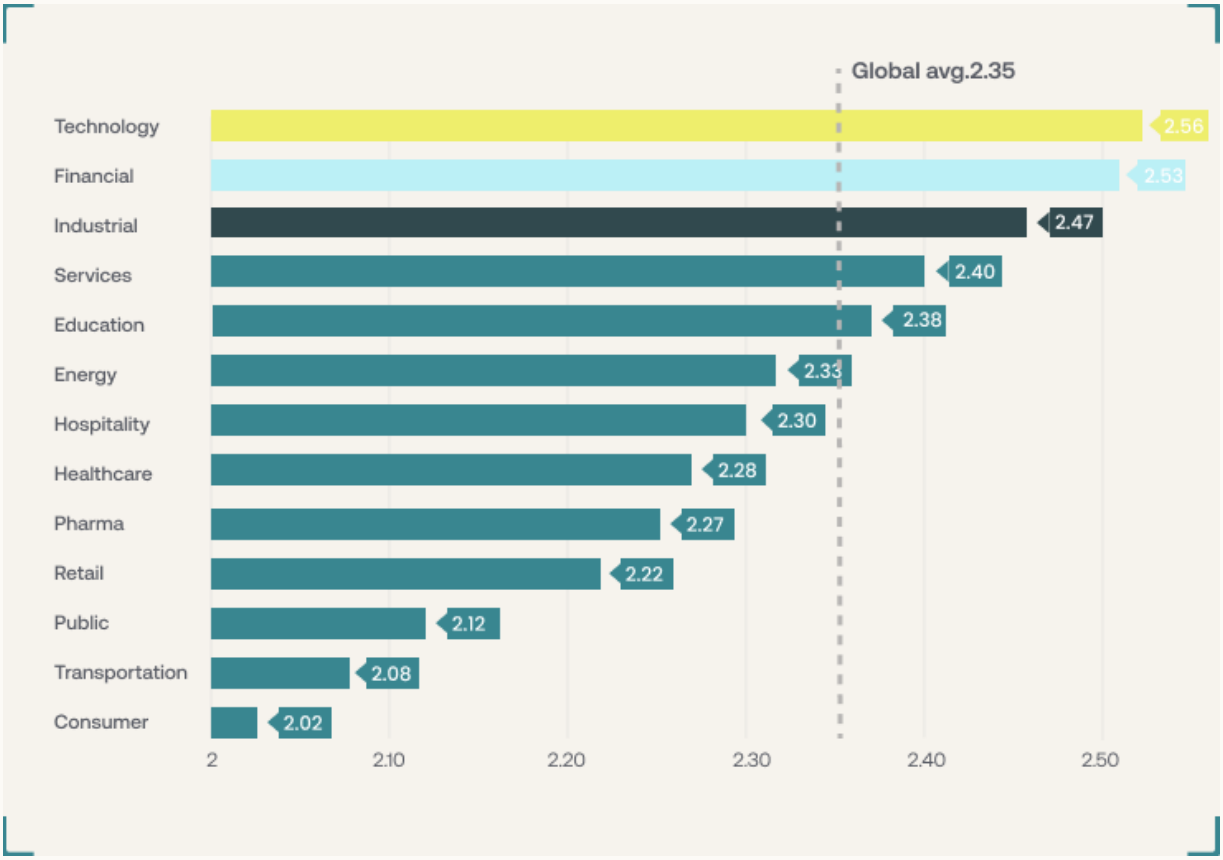

FEW INDUSTRIES LEAD ACROSS THE BOARD

Industry leaders vary by what you measure. Technology is the only sector strong across all four functions. Financial services governs and maps AI risk better than anyone, then falls to mid-pack on monitoring and management, it understands the risk better than it acts on it. Consumer and public-sector organizations lag everywhere.

AI MATURITY BY INDUSTRY AND FUNCTION

NIST AI RMF scores. Technology leads on Measure and Manage; Financial governs best but drops on execution.

	GOVERN	MAP	MEASURE	MANAGE
Technology	2.58	2.62	2.49	2.54
Industrial	2.41	2.50	2.47	2.47
Services	2.51	2.44	2.43	2.45
Education	2.43	2.38	2.35	2.40
Energy	2.48	2.29	2.26	2.29
Hospitality	2.29	2.28	2.34	2.33
Healthcare	2.39	2.33	2.29	2.33
Financial	2.76	2.68	2.40	2.27
Pharma	2.21	2.31	2.27	2.29
Retail	2.21	2.22	2.22	2.25
Public	2.18	2.13	2.06	2.15
Transportation	2.11	2.20	2.26	2.15
Consumer	2.11	2.11	1.93	2.02

AI Inherited Cybersecurity's Oldest Weakness

AI didn't create a new problem; it inherited cybersecurity's oldest one. In both frameworks, organizations score lowest exactly where awareness has to become action

That's why AI readiness can't be judged on its own. The AI maturity gap sits on top of a cybersecurity foundation that's still catching up, so to see where AI risk is really heading, we need to look at that foundation: the state of cybersecurity maturity itself.

2.18
NIST CSF
Protect

2.22
NIST AI RMF 1.0
Manage

Part Two

The Maturity Gaps Driving Real Risk

Maturity Is Improving - But Not Where It Counts

Cybersecurity maturity is generally improving across countries, industries, and company sizes, climbing from 2.19 to 2.35 in a year.

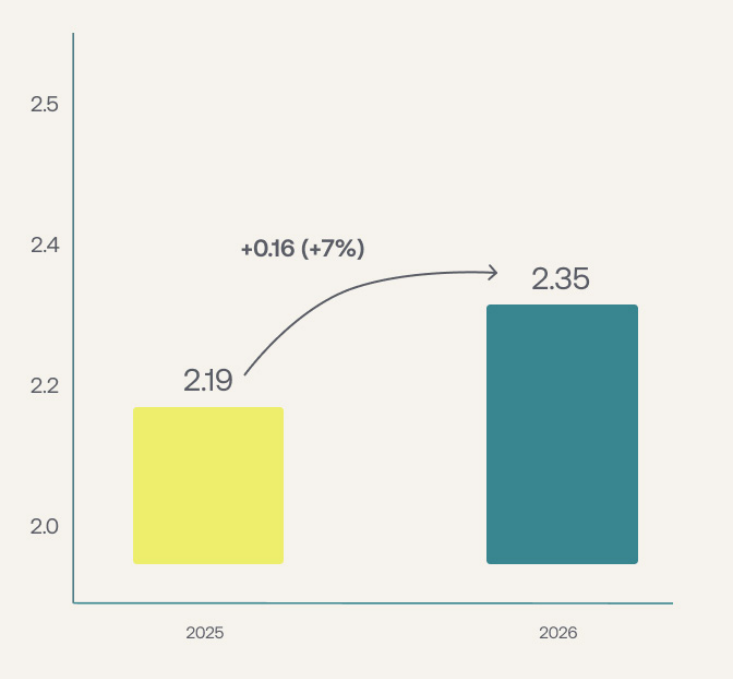

Real, broad-based progress, but organizations are stuck in the "Managed" band level, well short of "Defined" or “Optimized”. Organizations are more structured than before, just not yet consistent or operationalized at scale.

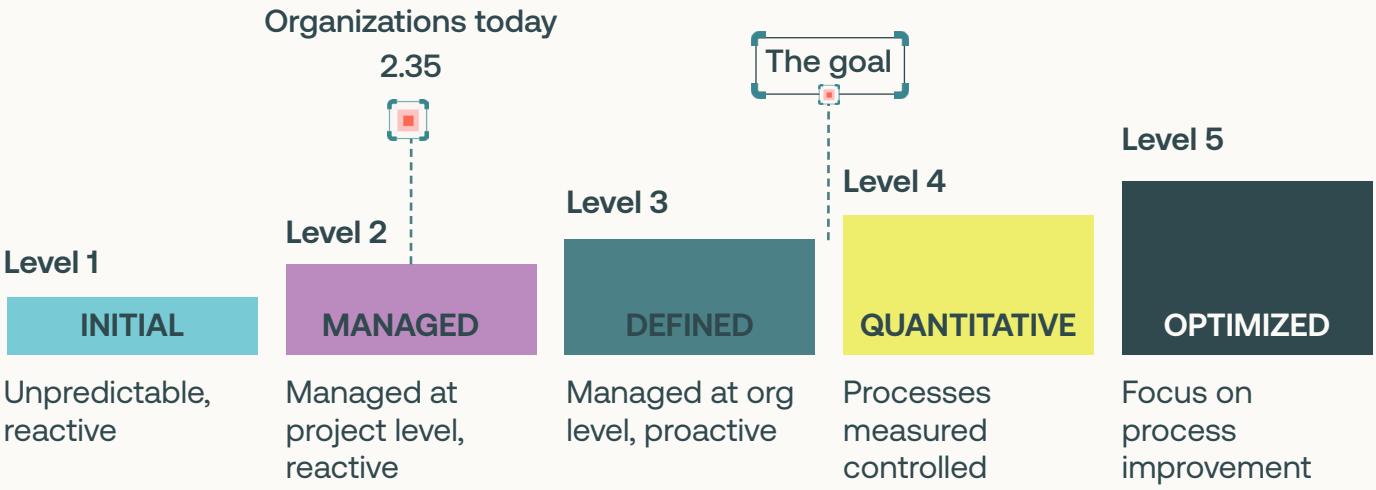

The boardroom shift is real, too. Cyber risk is now treated as business risk, debated alongside financial performance and operational resilience.

Attackers don't exploit board decks or policy documents. They exploit execution gaps, and that's exactly where the progress hasn't reached.

Prevention Is the Persistent Weakness

Years of investment have made organizations good at reacting to attacks, just not at preventing them. For the second year running, Protect is the lowest-scoring CSF function, and the gap is basic hygiene: the access controls, patching, and asset hardening that stop incidents before they start.

PREVENTION IS THE WEAKEST LINK

NIST CSF function scores. Strong at policy (Govern) and reaction (Respond), weak at prevention (Protect).

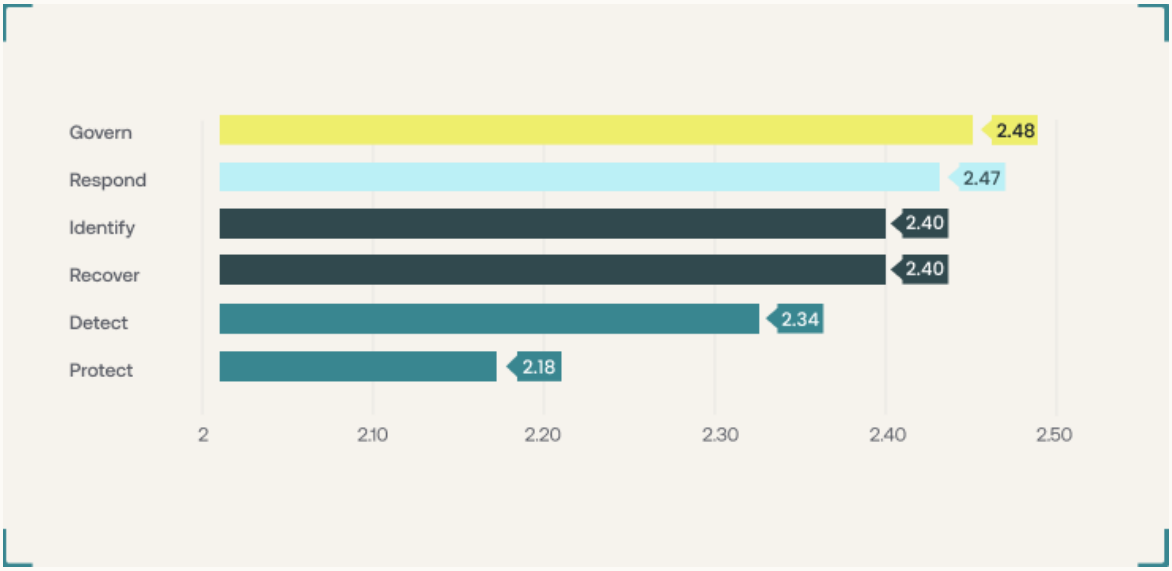

Organizations aren't disengaged, they're reactive. They score highest on the functions that bookend an incident, Govern (setting policy) and Respond (reacting once it happens), and lowest on Protect, the prevention in between. They can set the strategy and clean up afterward, but not stop the incident in the first place.

No one is immune; the gap persists across industry, geography, and budget. Even financial services, the strongest sector across almost every other function, lags on prevention. You can't respond your way out of exposure that grows faster than you can reduce it.

The Preventable Security Gaps Driving the Risk

Two weaknesses drive most of the exposure.

Basic Hygiene: the most common findings are outdated technologies, exposed admin interfaces, and missing web security headers, all known and fixable for years.

A Growing Cloud Blind Spot: insufficient monitoring, especially of cloud assets, leaves teams unable to see attacks as they unfold. As attackers move faster, that blindness becomes the difference between catching a vulnerability early and missing it entirely.

THE MOST COMMON FINDINGS ARE THE MOST BASIC

Top Cybersecurity finding by number of occurrences.
None are sophisticated; all are preventable.

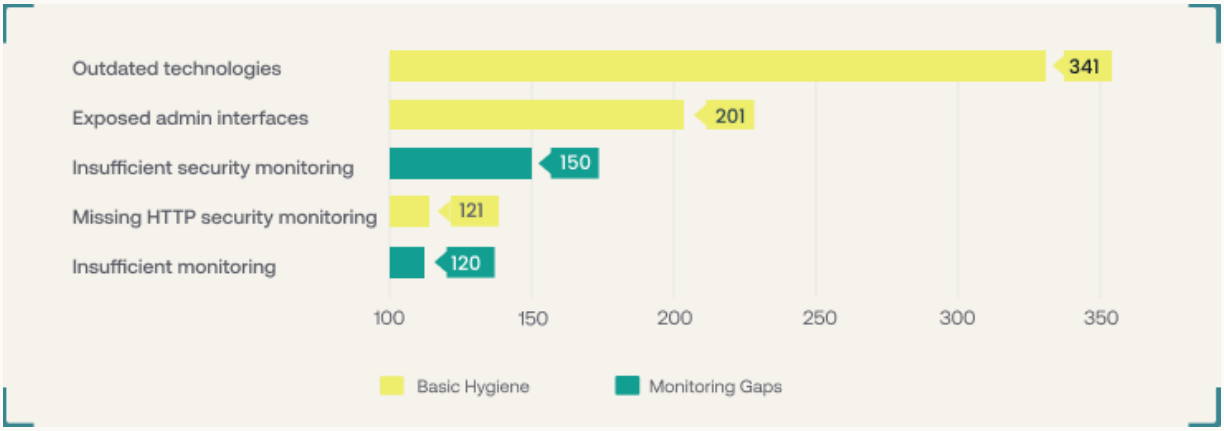

Maturity Scores Close Up

By Function, Industry, and Country

Maturity has improved globally from 2.19 to 2.35 in the past year. Across the six CSF functions, the awareness-to-action pattern holds: organizations are strongest at Govern and Respond, weakest at Protect.

Maturity Holds Up Unevenly

Maturity climbed to 2.35 this year, but progress is uneven across the six CSF functions. Organizations are strongest at Govern and Respond, and weakest at Protect, which falls furthest below the global average alongside Detect.

ABOVE AND BELOW THE GLOBAL AVERAGE

NIST CSF function scores.
Only Detect and Protect fall below the 2.35 average.

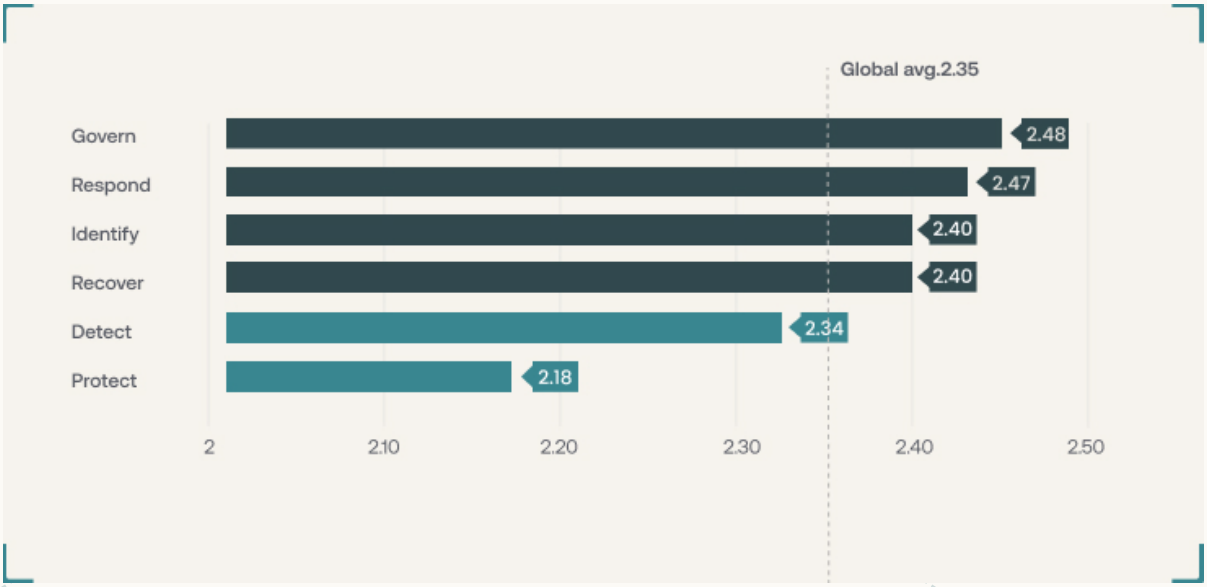

Industry Leaders by CSF Function

Financial services leads most functions, on the strength of regulation and investment. But it doesn't lead everywhere: the best incident-response scores belong to Healthcare and Hospitality, and on Protect, even the leading sector barely clears the floor. No industry is strong at prevention.

WHO LEADS EACH CSF FUNCTION

Top-scoring industry per function. Financial leads most, but the best Protect score barely clears the floor.

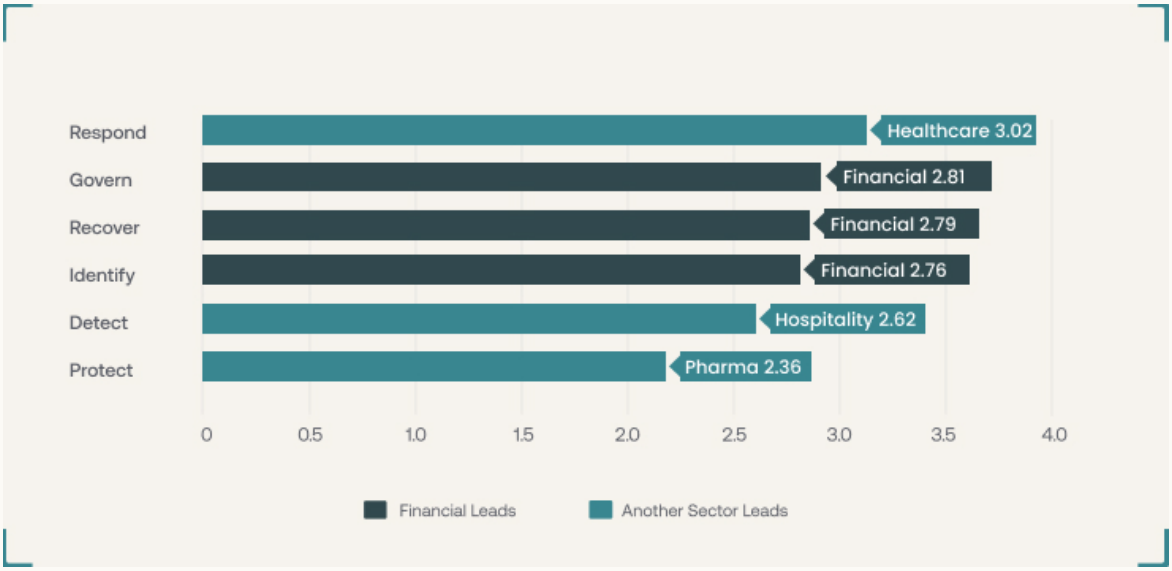

Where Countries Stand on Maturity

Maturity is rising everywhere, but no country has pulled clear. Only Switzerland, the United States, and Israel sit above the global average; every other country falls just below it. The spread is narrow and low, a reminder that even the leaders are still in the "Managed" band.

Switzerland's lead, built on regulatory investment paired with execution, shows what's possible when both move together.

CYBER MATURITY BY COUNTRY

Overall NIST CSF score, 2026. Only three countries sit above the global average

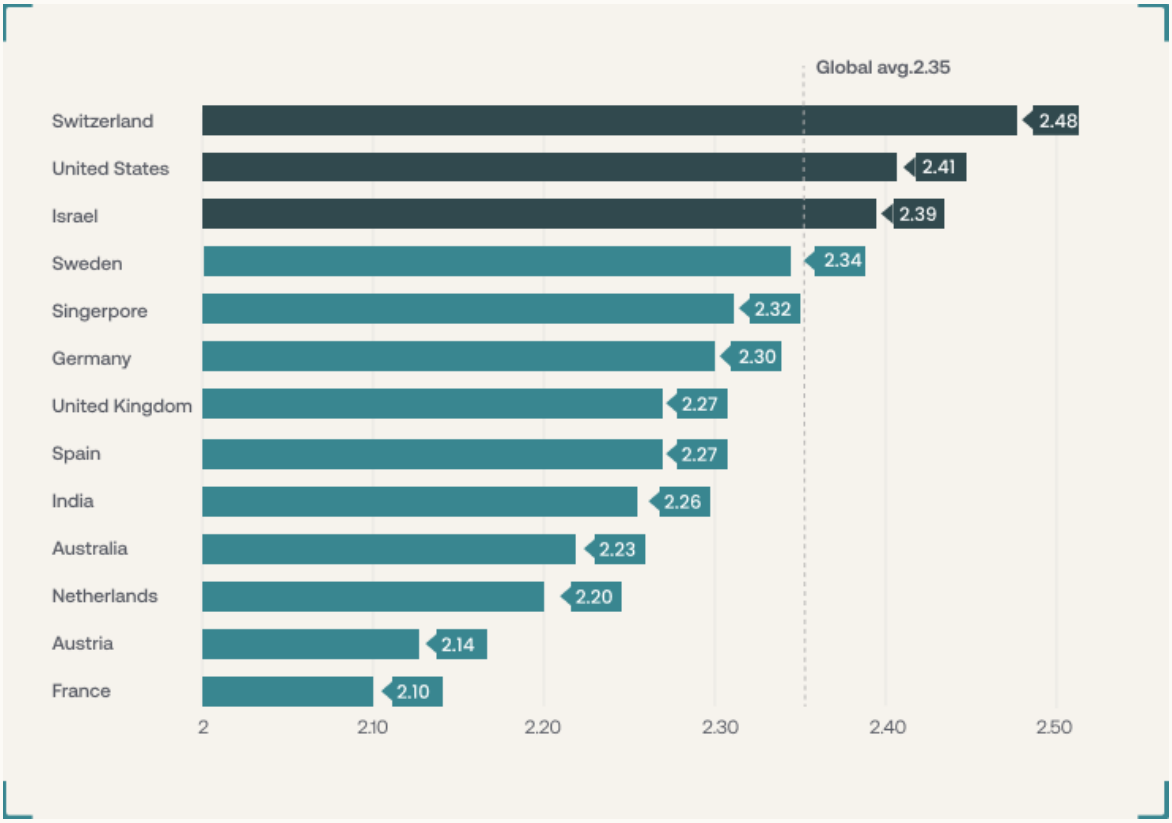

Part Three

What's Working, What Isn't, And What To Do

More Budget Doesn't Equal More Security

Spending more money isn't working.

Cybersecurity spending hit a record $212B in 2026, up 15% year over year, and 90% of organizations increased their budgets. Yet maturity hasn't moved out of the "Managed" band, and nearly a third of organizations feel less secure than a year ago.

Cybersecurity Spending
$212B
up 15% YoY - and 90% spent more

Yet

The Result
1 in 3
feel less secure than a year ago

The problem isn't how much organizations spend, it's where. Spending that misses critical functions buys activity, not resilience.

What Actually Works
Regulation Meets Execution

If money isn't the lever, regulation is, as long as it's backed by execution. The biggest maturity gains all followed regulatory deadlines, not bigger budgets.

Switzerland jumped +16% after new breach-reporting and resilience rules came into force, and the US, UK, and Spain each gained 8–9% behind their own.

MATURITY GAINS BY COUNTRY, 2025 TO 2026

Broad-based improvement across markets.
Switzerland led with a +16% jump.

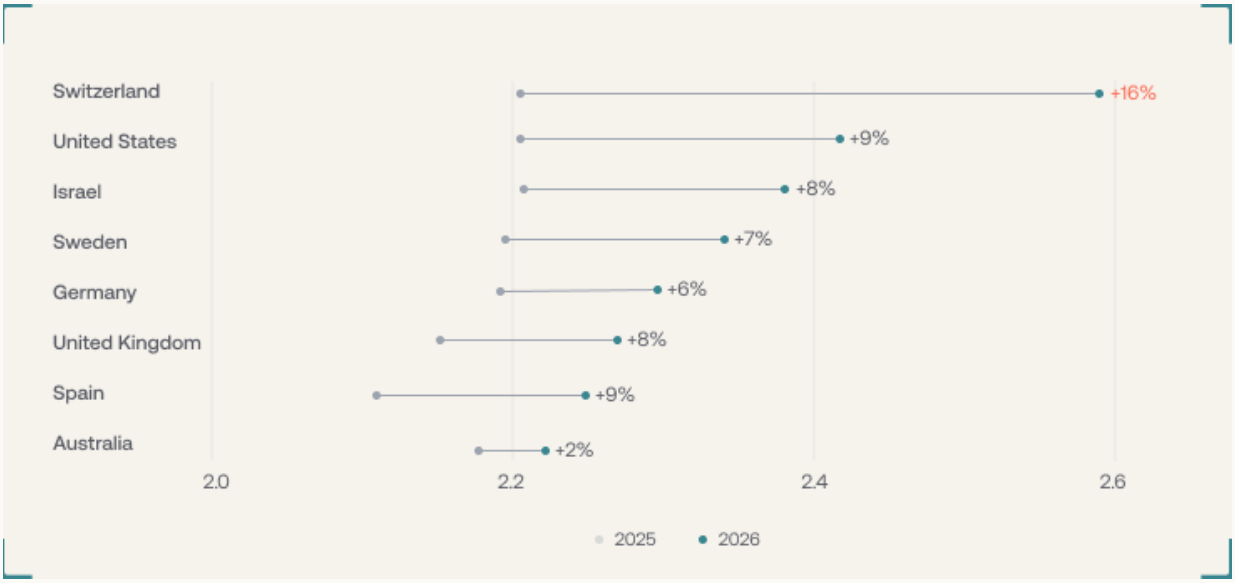

The pressure is rising. Full EU AI Act compliance is due by August 2026, with penalties up to €35M or 7% of global turnover, and SEC disclosure now extends to AI-related incidents, yet most industries still score below what those rules assume.

The Regulatory clock

August 2026

Full EU AI act compliance is due

Up to €35M or 7% of global turnover

in penalties yet most industries still score below what the rules assume.

But regulation isn't readiness on its own. France and Germany lead on AI policy yet sit mid-pack on maturity. The organizations pulling ahead treat compliance as a floor to build on, not a ceiling, and the divide shows: tightly regulated finance runs at just 5% Shadow AI exposure, against 71% in transportation and 62% in energy.

Cye Recommendations

How to Close the Gap

Knowing the gap isn't closing it. These steps move organizations from awareness to action, ordered from the foundation up.

1. **Map your AI before your next board meeting.** Inventory every AI tool, model, and service in use, including the shadow AI teams adopted on their own. You can't govern what you can't see, and this unlocks everything else.

2. **Give AI risk a named owner.** The CISO effect extends to AI: organizations with explicit executive ownership have far fewer unresolved AI findings. Assign the accountability, and the authority to act on it.

3. **Fix the basics first.** Protect is the lowest function two years running. Redirect spend to patch management, access controls, and security headers, the cheapest and most neglected path to maturity.

4. **Extend vendor risk to AI.** Insufficient vendor risk management is the single most common governance finding. Update assessments for AI-specific criteria, training data, model updates, data residency, and failure modes. Every third-party model is a dependency your program wasn't built for.

5 **Monitor AI before you scale it.** Most organizations can't yet detect AI-driven attacks. Cover model drift, pipeline integrity, prompt injection, and usage logging before deploying more. You can't respond to what you can't see.

6 **Drill the gap, don't document it.** Policy alone won't close the distance between governing and managing. Run tabletop exercises on real AI scenarios, prompt injection, data poisoning, a chatbot gone wrong. The teams that rehearse outperform the teams that only write policy.

7 **Use regulation as a floor, not a finish line.** The biggest maturity gains followed enforced deadlines, Switzerland's +16% chief among them, and the EU AI Act lands in August 2026. Treat compliance as the baseline to build past, and measure maturity continuously rather than once a year, so awareness keeps turning into action.

> **"The organizations that close 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

APPENDIX

Cye AI & Maturity Report 2026 — Data, Frameworks, and Methodology

The reference data behind the report's findings. All scores are on a 1–5 scale (see Scoring). Figures are drawn from the 2026 dataset.

A. About the Data

This year's report covers 21 countries, 16 industries, and 2,400+ NIST CSF function scores, up from 17 countries and 1,500 data points in 2025. Scores are drawn from thousands of cybersecurity maturity assessments conducted in Cye's exposure management platform. AI risk maturity was derived for the first time by mapping NIST CSF 2.0 function scores to the four core functions of the NIST AI RMF 1.0.

SCORING SYSTEM (CMM)

Scores follow the Capability Maturity Model (CMM), on a scale of 1 to 5, where 1 is the least mature (most vulnerable) and 5 is the most mature.

FUNCTION	SCORE	DESCRIPTION
Initial	1	Ad hoc, reactive
Managed	2	Basic processes in place
Defined	3	Standardized, documented
Quantitative	4	Measured and controlled
Optimized	5	Continuous improvement

B. The Frameworks

NIST CYBERSECURITY FRAMEWORK (CSF) 2.0 — SIX FUNCTIONS

FUNCTION	FOCUS
Govern	Strategy, risk tolerance, policy, and oversight
Identify	Asset inventory and risk assessment
Protect	Safeguards, access controls, and preventive measures
Detect	Monitoring and anomaly detection
Respond	Incident response, analysis, and communication
Recover	Restoration and business continuity

NIST AI RISK MANAGEMENT FRAMEWORK (AI RMF) 1.0 — FOUR FUNCTIONS

FUNCTION	FOCUS
Govern	Leadership & oversight; policies & procedures; culture
Map	System inventory; stakeholder engagement; context analysis
Measure	Risk assessment; performance monitoring; metrics & tools
Manage	Risk mitigation; incident response; continuous improvement

C. AI Risk Maturity — Data Tables

GLOBAL AI RMF FUNCTION SCORES

FUNCTION	SCORE
Govern	2.46
Map	2.40
Measure	2.33
Manage	2.22

AI GOVERN BY COUNTRY

COUNTRY	AI GOVERN
Switzerland	2.56
United States	2.40
Israel	2.38
Singapore	2.32
Sweden	2.31
United Kingdom	2.29
Germany	2.28
Spain	2.26
Denmark	2.26
India	2.25

Country-level AI data was available for the Govern function only.

SHADOW AI EXPOSURE BY INDUSTRY

% of organizations using AI beyond formal governance (below-median AI Govern with active AI-related findings).

INDUSTRY	EXPOSURE
Transportation	71%
Energy	62%
Industrial	42%
Healthcare	38%
Technology	14%
Services	8%
Financial	5%

Every organization assessed had findings tied directly to AI systems. Three categories dominate: AI Infrastructure Misconfiguration (48), AI-Specific Attacks (3), and AI Governance Gaps (2).

AI-RELATED FINDINGS: KEY ACCESS NOT DISABLED	COUNT
Cloud AI Services: key access not disabled	9
Cloud AI Services: network access unrestricted	8
Cloud ML Workspaces: public network access enabled	5
Cloud ML Workspaces: resource logs not enabled	5
Cloud ML Computes not in a virtual network	4
AI engine vulnerable to prompt injection	2
Unrestricted use of GenAI (data privacy)	1
Informal accountability in AI risk management	1

D. Cybersecurity Maturity — Data Tables

CSF 2.0 FUNCTION AVERAGES

COUNTRY	AI GOVERN
Govern	2.48
Respond	2.47
Recover	2.42
Identify	2.41
Detect	2.34
Protect	2.18

CSF FUNCTION DEEP DIVES — TOP INDUSTRIES AND COUNTRIES

Govern

TOP INDUSTRIES	SCORE	TOP COUNTRIES	SCORE
Financial	2.81	Israel	2.96
Technology	2.74	Switzerland	2.75
Hospitality	2.74	Denmark	2.71
Energy	2.72	Singapore	2.69
Education	2.69	USA	2.65

Identify

TOP INDUSTRIES	SCORE	TOP COUNTRIES	SCORE
Financial	2.76	Switzerland	2.81
Technology	2.67	Israel	2.79
Healthcare	2.61	Singapore	2.62
Industrial	2.57	USA	2.59
Pharma	2.52	Denmark	2.55

Protect

TOP INDUSTRIES	SCORE	TOP COUNTRIES	SCORE
Pharma	2.36	Switzerland	2.53
Financial	2.35	Singapore	2.44
Technology	2.34	United Kingdom	2.43
Hospitality	2.29	USA	2.31
Services	2.28	Denmark	2.27

Detect

TOP INDUSTRIES	SCORE	TOP COUNTRIES	SCORE
Hospitality	2.62	Switzerland	2.78
Financial	2.54	Israel	2.67
Industrial	2.52	Singapore	2.59
Services	2.49	Denmark	2.51
Technology	2.45	USA	2.51

Respond

TOP INDUSTRIES	SCORE	TOP COUNTRIES	SCORE
Healthcare	3.02	Israel	3.05
Hospitality	3.00	Switzerland	3.04
Financial	2.96	Singapore	2.94
Technology	2.96	Denmark	2.91
Industrial	2.85	USA	2.85

Recover

TOP INDUSTRIES	SCORE	TOP COUNTRIES	SCORE
Financial	2.79	Switzerland	2.73
Pharma	2.66	Netherlands	2.71
Hospitality	2.64	Spain	2.65
Entertainment	2.62	Denmark	2.64
Education	2.57	United Kingdom	2.59

OVERALL CYBER MATURITY BY COUNTRY (2026)

COUNTRY	SCORE	COUNTRY	SCORE
Switzerland	2.58	Denmark	2.27
United States	2.41	India	2.26
Israel	2.39	Australia	2.23
Sweden	2.34	Netherlands	2.20
Singapore	2.32	Austria	2.14
Germany	2.32	France	2.10
United Kingdom	2.30		
Spain	2.27		

MATURITY BY COUNTRY — 2025 VS 2026

COUNTRY	2025	2026	CHANGE
Switzerland	2.22	2.58	+16%
United States	2.21	2.41	+9%
Spain	2.09	2.27	+9%
United Kingdom	2.13	2.30	+8%
Israel	2.21	2.39	+8%
Sweden	2.19	2.34	+7%
Germany	2.18	2.32	+6%
Australia	2.18	2.23	+2%
Norway*	2.58	—	Retained
Japan*	2.45	—	Retained
Mexico*	1.73	—	Retained

* Assessed in the 2025 cycle but not 2026; 2025 scores retained for continuity.

MOST FREQUENT FINDINGS BY CSF FUNCTION

FUNCTION	TOP FINDINGS OCCURRENCES
Govern	Insufficient Vendor Risk Management (91); Weak Password Policy (71); Password Reuse (58)
Identify	Outdated Technologies (341); Admin Interfaces Exposed (201); No HSTS (121)
Protect	Outdated Technologies (681); Missing HTTP Headers (578); Information Disclosure (540)
Detect	Insufficient Security Monitoring (150); Insufficient Cloud Monitoring (120); Corporate Network Monitoring (83)
Respond	Insufficient IR Procedures (54); Insufficient IR – Financial (16); Insufficient IR – OT (16)
Recover	Key Vaults Missing Soft Delete (10); Missing Business Continuity Plan (8); No Disaster Recovery Plan (7)

E. Methodology

Each assessment was conducted in Cye's exposure management platform and scored against the six NIST CSF 2.0 functions, with detailed findings, severity levels, affected assets, and recommended mitigations. Results were parsed automatically; scores and top findings were extracted; and the data was aggregated by industry, organization size, and country while preserving client confidentiality. AI risk maturity scores were derived by mapping CSF 2.0 function scores to the four NIST AI RMF 1.0 functions, with AI-related findings identified through keyword analysis and NIST subcategory classification. Findings reflect the 2026 assessment cycle.