HACKTHEBOX

The Global Cyber Skills Benchmark 2025

Where 795+ cyber teams were tested—and what their performance revealed

Table of contents

Executive summary

Cybersecurity teams around the world are getting better. But not fast enough—and in some industries, not where it matters most.

This year's Global Cyber Skills Benchmark reveals a growing divide between elite performers and the average team. Foundational domains like Secure Coding (18.7%), Cloud (21.3%), and Web Security (21.1%) continue to be major weak spots, even in high-risk industries.

Geography offered a few surprises. Countries like Japan (52.1%) and Vietnam (40.2%) outpaced traditional leaders like the US (30%) and the UK (31.6%). But then, some of the top 10 teams across the whole CTF were US-based—proving that standout talent exists everywhere, even when national averages suggest otherwise.

AI and Machine Learning (ML) emerged as new challenge categories this year, with solve rates of 37% and 30.1%, respectively. Despite the hype, nearly 93% of teams said AI tools were not essential to their success, underscoring that human expertise still drives results.

Executive summary

Industry shake-ups were just as telling. Business Services led all sectors with a 43.9% average solve rate, while Government (27%) and Education (20.4%) fell to the bottom. The introduction of emerging threats—like AI prompt injection and smart contract exploits—reshuffled the leaderboard and exposed blind spots in training.

The takeaway is clear: certifications may check compliance boxes, but attackers don't care about credentials. Real resilience is performance-based. If your team can't handle cloud misconfigs, lateral movement, or adversarial AI, policy alone won't protect you.

This report is designed to move organizations from insight to action. We unpack what top teams are doing differently, identify where the biggest skill gaps remain, and explore how frameworks like CTEM—Continuous Threat Exposure Management—can help close those gaps with real-world, data-backed strategies.

Let's get into it.

796
Security teams

4,549
Professionals

40+
Challenges

$50,000+
Prize pool

3 of 10 top performing teams were from **US-based** companies

Benchmarking global threat response capabilities

What ‘good’ looks like: Inside high-performing teams

While global solve rates improved across most categories compared to 2024, a handful of elite teams dramatically outpaced the rest. Understanding what these top performers have in common provides a playbook for organizations looking to close critical skill gaps.

Looking at the highest-performing teams—those scoring **above 90% across key domains like Coding, Forensics, and Cloud**—a pattern emerges. These teams don’t just train harder. They train smarter, with deliberate, real-world alignment.

What sets top performers apart?

- Frequent benchmarking across live threat domains, not static frameworks
- Secure Development Lifecycle (SDLC) integration across projects
- Red teaming and lateral movement testing to expose critical weaknesses
- Real-time exposure management mapped to CTEM (Continuous Threat Exposure Management)
- Strategic use of AI for support tasks (e.g., syntax generation, decoding), never as a crutch

These aren’t hypothetical best practices—they’re active behaviors that can be inferred from performance across the teams that dominated this year’s competition.

Where top talent comes from

Most of the top-performing teams came from IT services companies—a sector where cyber capability is core business.

But not exclusively: finance and technology firms also cracked the leaderboard, showing that in-house talent can thrive when given the right environment and support.

Characteristics of top performing teams

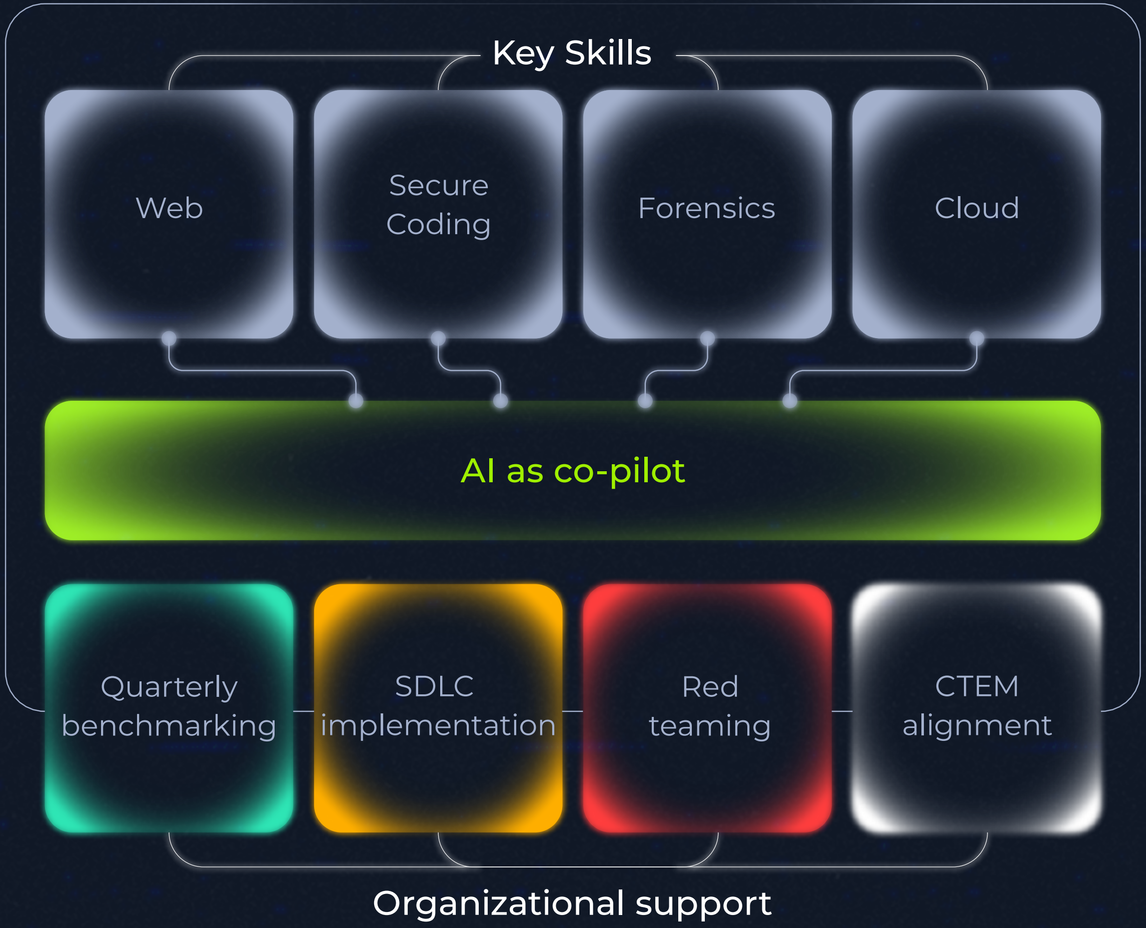

Not just the usual suspects

High performance isn't tied to geography. This year's results showed countries like **Japan (52.1%), Vietnam (40.2%),** and **Switzerland (40.2%)** outpace traditional cyber leaders like the **UK (31.6%)** and **US (30%)** at the country level.

But, when we take a closer look at the details, we see a different story come to the surface: **three of the top 10 teams in the entire benchmark were US-based**, proving that elite talent exists even where average performance lags.

The skills gap persists

Despite general improvements across categories, some of the most mission-critical skill sets remain underdeveloped

- **Secure Coding:**
  Just 18.7% of challenges solved. This low performance suggests a continued lack of Secure Development Lifecycle (SDLC) implementation across industries.
- **Web Security:**
  With only 21.1% of challenges solved, web app misconfigurations remain an overlooked threat.
- **Cloud Security:**
  At 21.3%, cloud challenge performance reflects an ongoing struggle with securing modern infrastructure.

These aren't fringe skills—they're the foundation of modern digital environments. And yet, they seem to be consistently falling by the wayside.

By contrast, **Coding (53.6%) and Forensics (48.2%)** showed the strongest year-over-year growth, suggesting that teams are getting better at incident response and automation. Meanwhile, **OSINT debuted as the top-performing category (65.7%)**, hinting at growing organizational maturity in reconnaissance and threat intelligence.

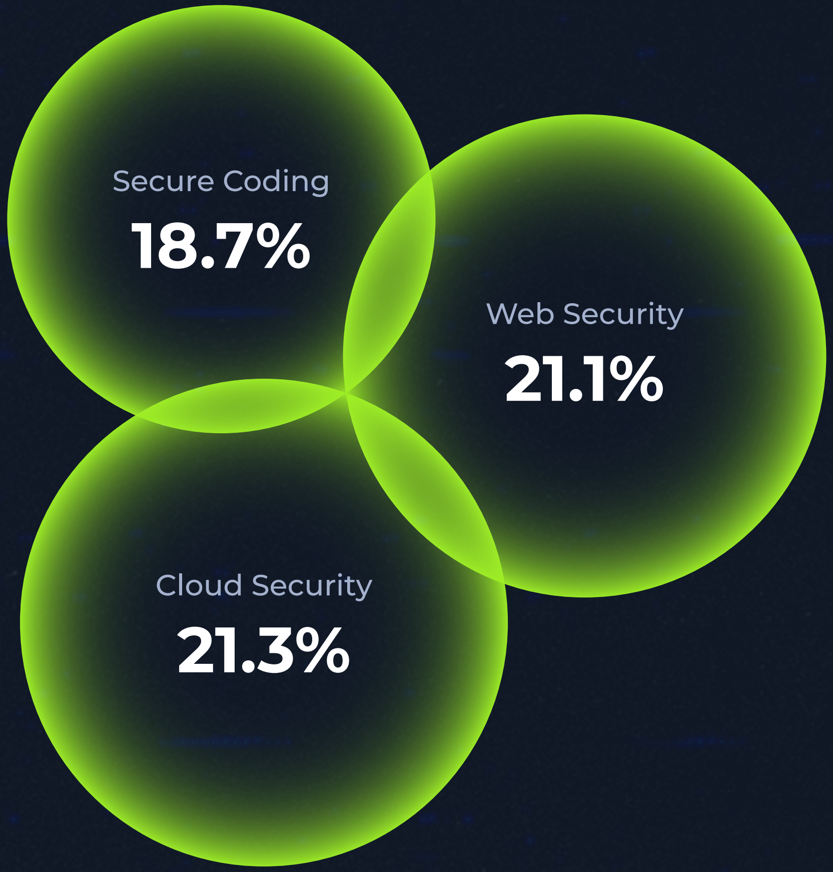

AI is a co-pilot, not a crutch—for now

This year introduced AI and Machine Learning as challenge categories, with solve rates of 37% and 30.1% respectively—impressive for emerging domains. But it wasn't just the challenge data that stood out. **Nearly half (44%) of teams reported using AI tools during the CTF**, but only 7.5% said those tools were crucial to solving challenges.

In practice, AI was used to support—not replace—human reasoning: helping with decoding, syntax generation, and concept clarification. It's a sign that AI is a useful assistant, but not a shortcut to success.

However, this divide also reveals a deeper risk. **Secure Coding remained one of the weakest categories (18.7% solve rate, 0% in Education)**, raising red flags about overreliance on AI in development workflows. The speed AI provides can amplify underlying weaknesses, introducing subtle, hard-to-detect vulnerabilities.

As AI becomes embedded in day-to-day tasks, Secure Coding must evolve from an afterthought to a core skill. Otherwise, AI won't be a competitive edge. It'll become a liability.

Training for the age of AI

AI isn't coming—it's already here. But while threat actors are racing to weaponize it, many defenders are still learning how to integrate it meaningfully into their workflows. The next frontier of cyber defense won't just rely on technical expertise or threat intelligence.

It will require the ability to work with intelligent systems, whether that involves detecting, responding to, or outmaneuvering adversaries in real-time.

As AI becomes embedded in security tooling, workflows, and even adversary techniques, upskilling must also evolve. It's not enough to learn how AI works. Teams need realistic hands-on experience in how it behaves under pressure; how it can be guided, manipulated, or even exploited. That kind of fluency only comes from disciplined practice.

That's why we've introduced a new layer to our CTF environment: one that lets teams train not just on AI topics, but with AI tools, agents, and protocols.

A new kind of cyber range

To support this, Hack The Box has launched the **Model Context Protocol (MCP)**—a framework that allows the integration of AI agents and custom tooling directly into CTF challenges. Think of it as an API for adversarial simulation: teams can build, invoke, and chain AI-native workflows, mimicking how real-world security operations might evolve in AI-augmented environments.

From SIEM integrations and autonomous red teaming to agent-driven scripting, MCP turns traditional training into a hybrid human-AI exercise.

Because the future of cybersecurity won't be human or machine. It'll be both. And your team needs to prepare for that reality now.

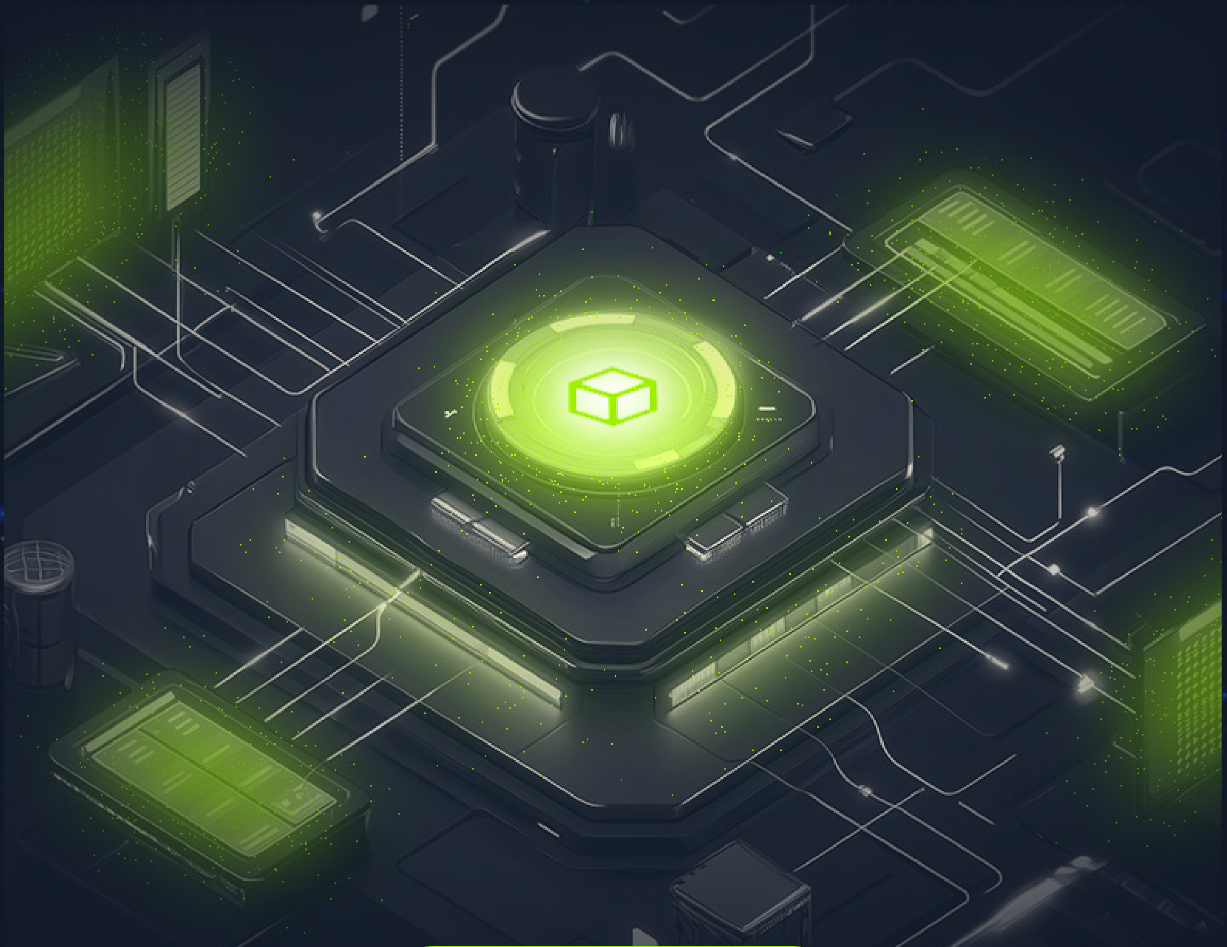

See HTB MCP in action

A new kind of cyber range

Detection ≠ Defense

Strong visibility doesn't guarantee resilience. Some sectors excel at spotting threats but fail to neutralize them. Take Retail: an industry-leading 79.7% solve rate in OSINT, but just 20.3% in Web and 20.8% in Secure Coding.

This disconnect is common—and dangerous. High detection without the skills to act breeds overconfidence and leaves critical systems exposed.

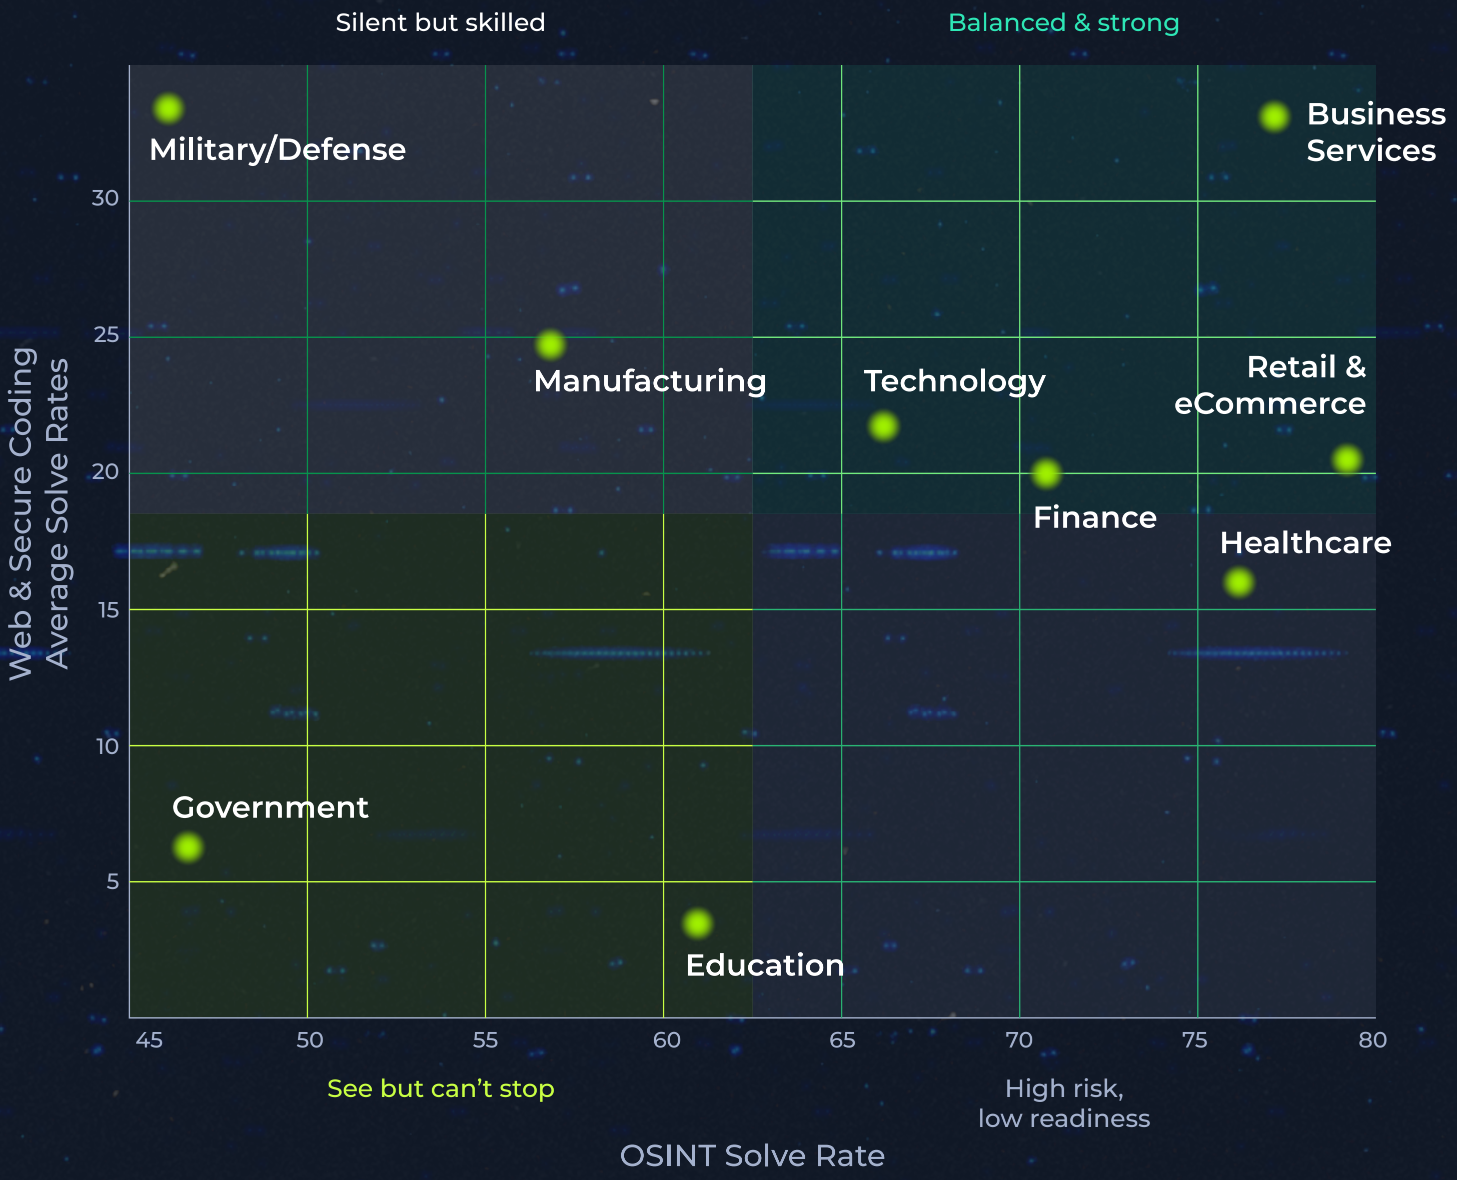

Who's winning the global cyber skills race?

In something of a surprise twist, the US appeared slightly lower in the country-level rankings—but when we look at the individual company performance, we see a different story.

Three of the top 10 teams, including one in third place, entered the CTF through US-based companies. This highlights that strong performers are there, even if they're not reflected in the aggregate statistics for that country.

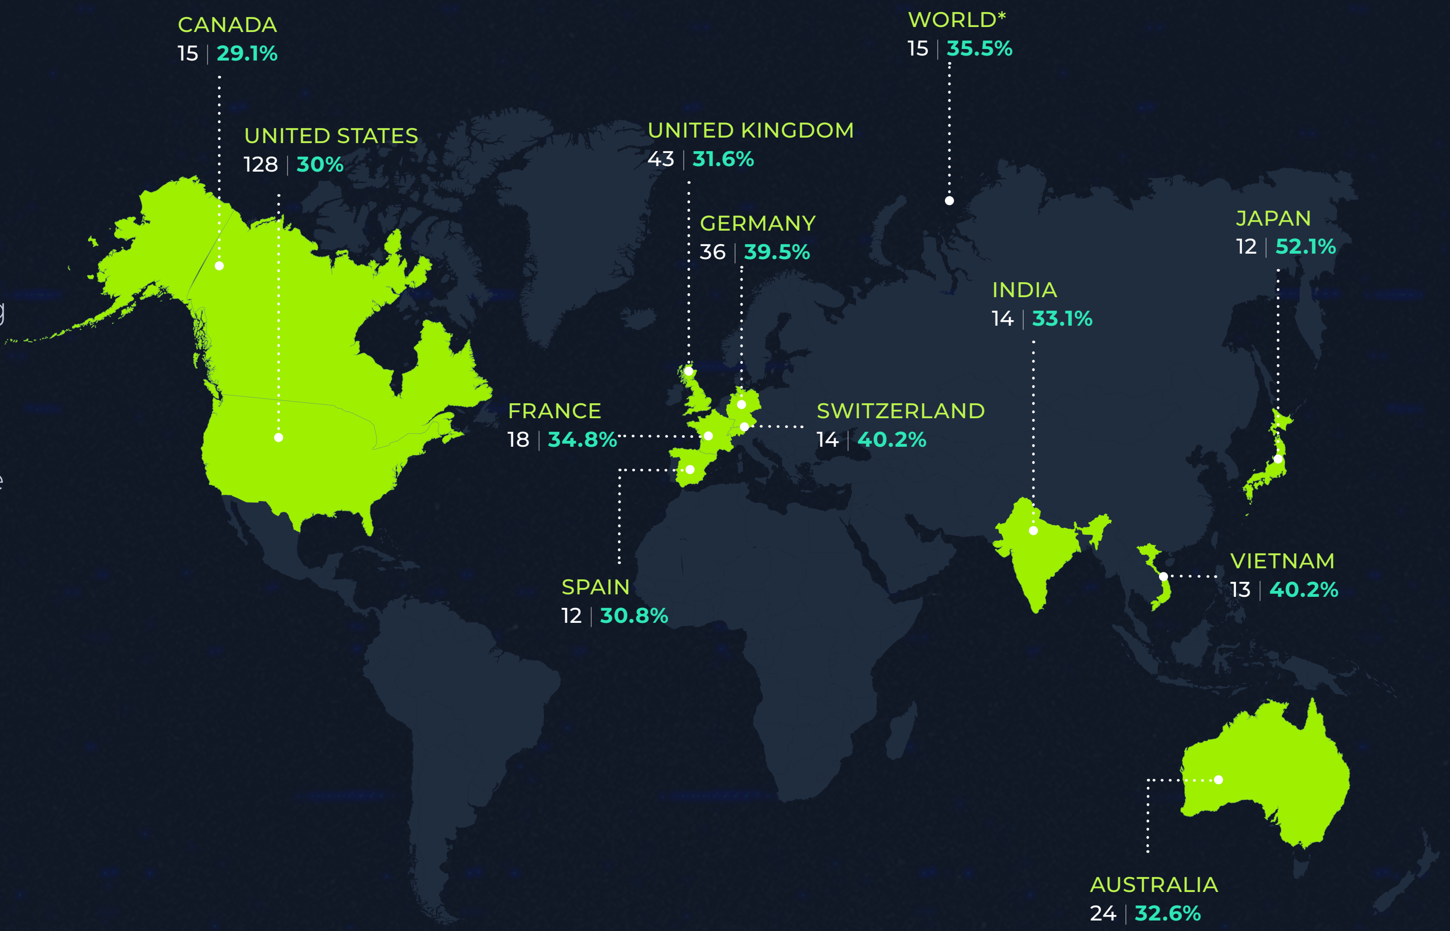

⬢ Number of teams

⬢ Average solves (All categories)

*This field denotes companies that participated but did not disclose their specific location.

Top 10 companies and where they're from

The best-performing teams mostly came from cybersecurity services companies—but a few standouts from tech and finance proved that in-house talent can also rise to the top.

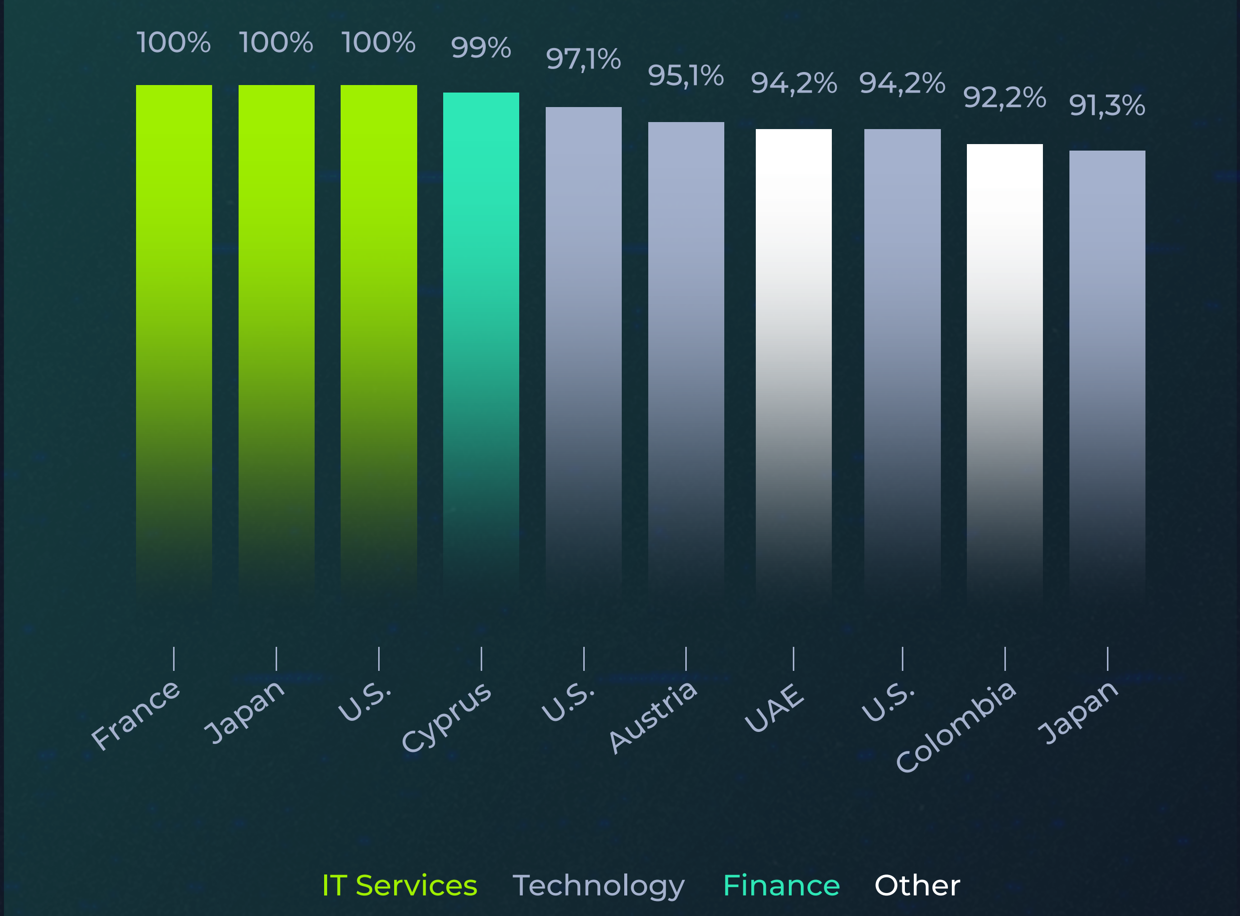

The leaderboard by industry

Industries with strong baseline IT maturity performed best. **Business Services topped the chart (43.9%)**, a dramatic turnaround from last year. **Military/Defense followed closely (41.4%)**, showing high competence in Coding (72%) and Forensics (62.8%).

Finance averaged a respectable 37.6% of solves across all challenges, excelling in areas like OSINT (71%), Forensics (54.6%), and Coding (51.4%) but falling behind in Cloud (22.9%) and Secure Coding (20.8%).

On the flip side, **Education (20.4%) and Government (27%)** saw steep declines. That's troubling given the critical importance of cyber resilience in public infrastructure. Education, in particular, had the lowest score in **Secure Coding (0%)**—a major red flag.

Healthcare (27.5%), **Retail and eCommerce (33.2%)**, and **Energy and Utilities (31%)** industries all hovered around the lower third, with similar weaknesses in Fullpwn, Web, and Secure Coding.

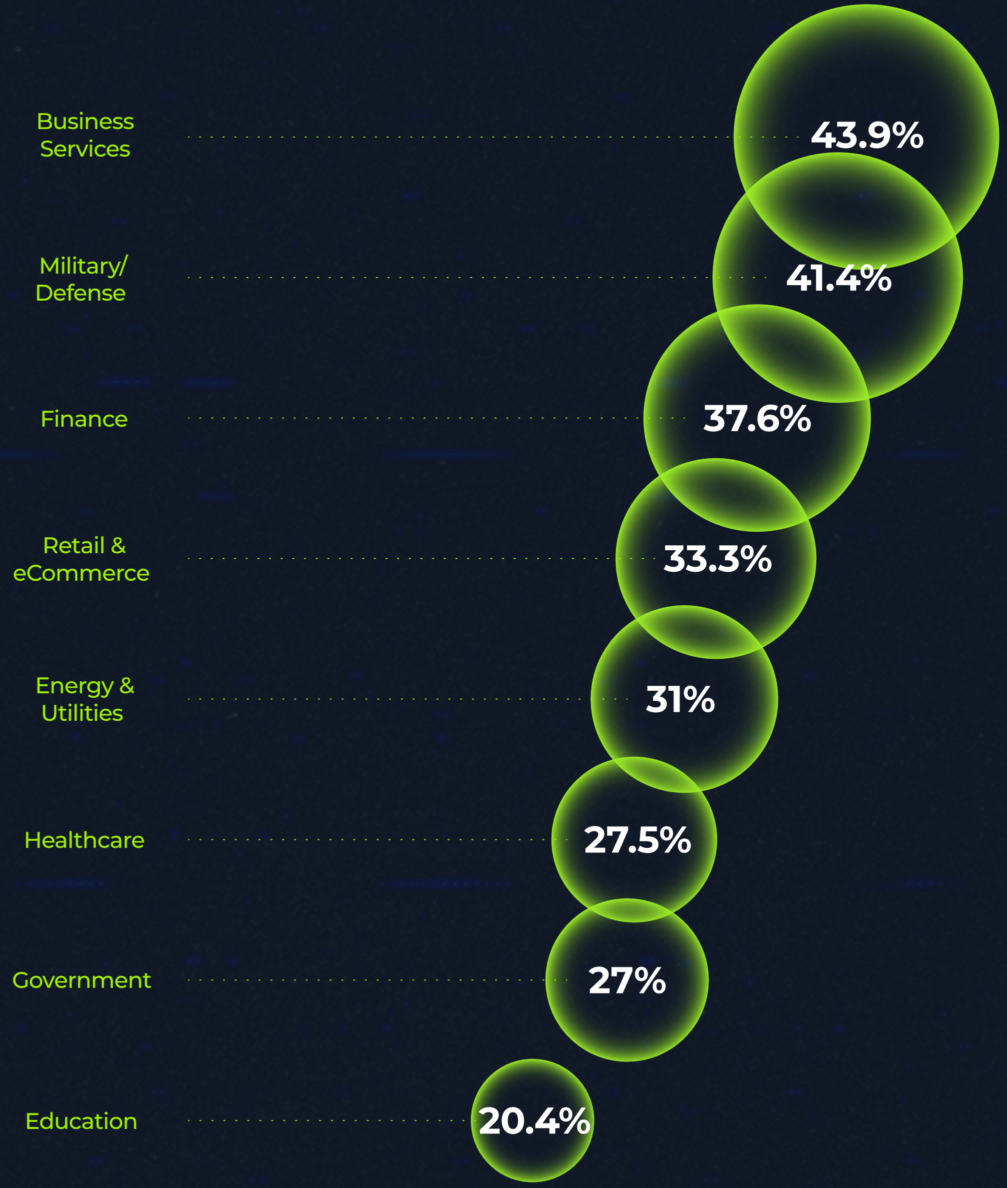

The leaderboard by industry

This year's outliers

- Finance underperformed in Blockchain challenges (10.1%), despite being one of the sectors most likely to explore Web3.
- Technology teams did well in AI/ML and scored consistently above average across categories.
- Manufacturing ranked highest in Pwn (10.9%) and Reversing (29.1%), hinting at niche strengths in embedded systems.

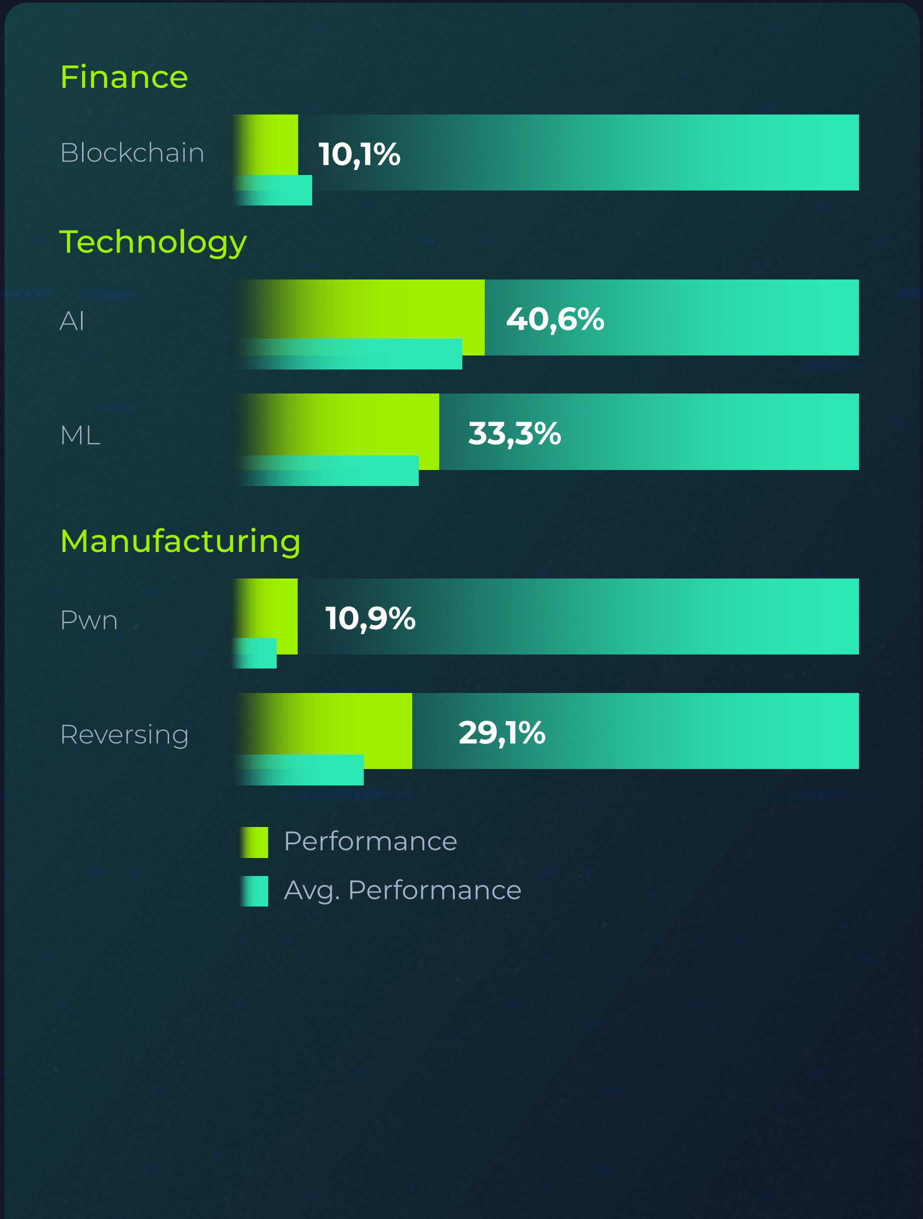

Industry spotlights

Retail and eCommerce

Retail and eCommerce organizations have become increasingly attractive targets for cybercriminals. High transaction volumes, complex supply chains, and rich customer data make them vulnerable to ransomware, credential stuffing, and third-party breaches. From point-of-sale malware to attacks on payment platforms and loyalty programs, the retail threat landscape has become relentless.

And yet, this year's benchmark shows the sector still has some work to do. Despite solid performance in **OSINT (79.7%)** and **AI challenges (45%)**, retail teams scored **below average in key defensive domains like Secure Coding (20.8%)**, Web (20.3%), and Cloud (22.5%)—the very areas attackers routinely exploit in retail-specific breaches.

Threat groups like Scattered Spider have moved beyond tech and hospitality, targeting retail-adjacent brands through cloud misconfigurations, web vulnerabilities, and identity compromise.

Earlier breaches at the likes of Shopify merchants and Fast Retailing Co. (Uniqlo) exposed similar weaknesses: insecure payment flows, web app flaws, and supply chain risks. These are the same gaps highlighted in this year's benchmark, proving they're not just training issues, but real attack vectors.

For retail and eCommerce, the message is clear: close the gap before attackers exploit it.

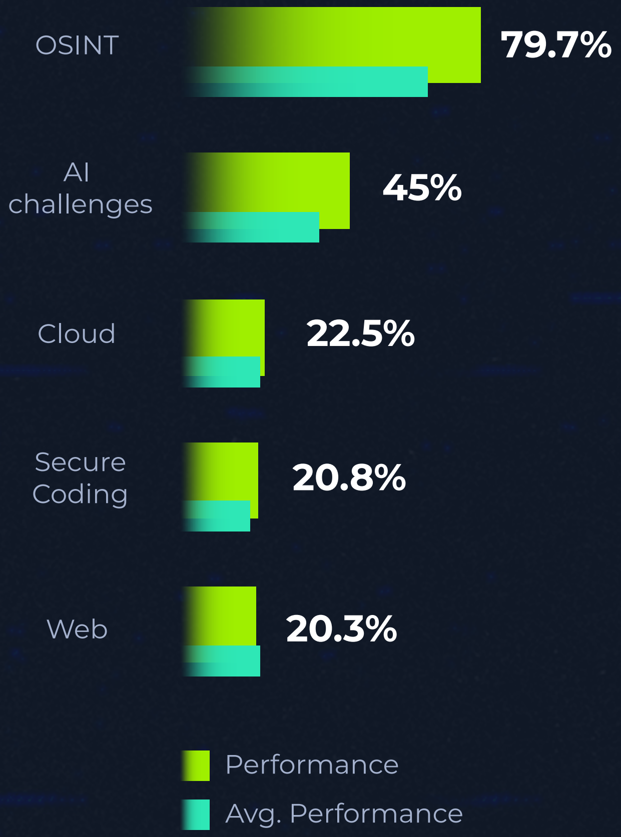

Retail and eCommerce

> Minimizing the impact of cyberattacks requires meaningful investment in people, tools, and training—alongside a culture that supports security-driven decision-making.
>
> Attack Surface Management, IT hygiene, and proactive defense must be ongoing disciplines. Teams need to be battle-ready through regular purple team exercises, real-world threat simulations, and performance evaluations.
>
> Simulating cyber incidents in end-to-end scenarios helps stress-test response plans, break down silos, and build faster, more resilient incident response
>
> **Andrew Morris**
> Head of Defensive Content Engineering @ Hack The Box

Key takeaways

- Retailers should prioritize simulated supply chain attacks and web app vulnerabilities in training programs.
- Emphasize secure development practices to protect customer data and payment flows.
- Leverage existing strengths in OSINT and forensics for faster breach detection and response.
- Given AI strength, there’s potential to integrate automated monitoring or threat detection tooling.

Government and Military/Defense

Cybersecurity is a matter of national security—but the skills data suggests a widening gap. While **Military/Defense teams scored highly in foundational areas like Coding (72%) and Forensics (62.8%)**, broader Government sector performance lagged in emerging domains.

Government teams scored just **17.8% in AI, 5.6% in Blockchain**, and **8.9% in Cryptography**. In short, domains that are rapidly becoming central to both cyber offense and defense. These results suggest a critical vulnerability: while leadership embraces innovation, the operational workforce often lacks the hands-on capability to secure it.

These are precisely the areas adversaries are targeting—and that defense leaders like USCYBERCOM are racing to secure. In 2024–2025, the command launched a sweeping AI roadmap with over 100 initiatives across cyber operations and threat disruption.

But here's the disconnect: while leadership is embracing AI, many government teams lack the hands-on skills to use it securely. That has real consequences; as adversaries deploy GenAI for phishing, impersonation, and payload generation, the public sector can't afford to treat AI as a black box. Without the hands-on skills to understand and defend these systems, policy outpaces capability, and that's a national security risk.

Low engagement with advanced topics like GenAI, Blockchain, and cryptography points to strategic vulnerabilities—especially as attackers lean into automation, prompt injection, and adversarial ML.

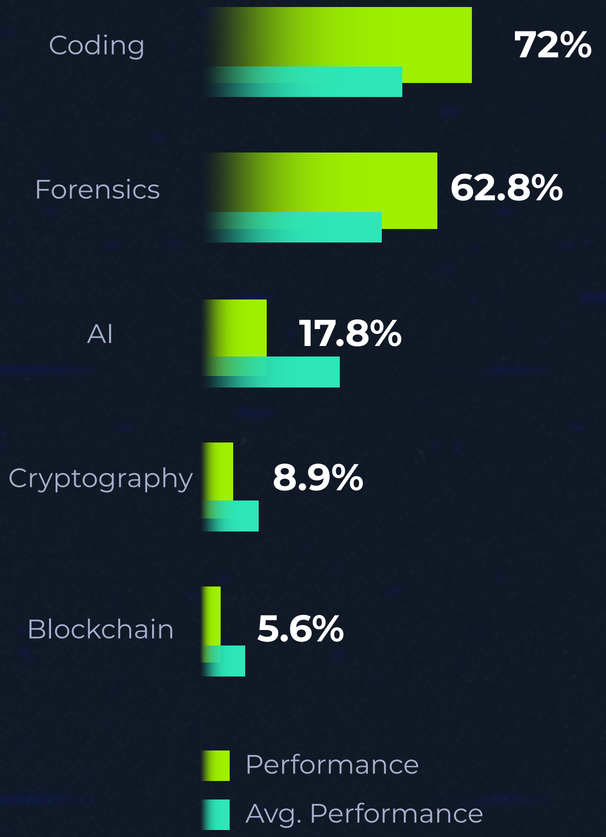

Government and Military/Defense

Key takeaways

- Foundational skills are solid. Military/Defense teams performed exceptionally well in Coding (72%) and Forensics (62.8%), showing strength in traditional cybersecurity operations.
- Emerging tech gaps remain. Government teams struggled in critical future-facing areas: AI (17.8%), Blockchain (5.6%), Cryptography (8.9%).
- Low AI engagement is a liability. Despite USCYBERCOM's new roadmap to integrate AI across cyber operations, the hands-on skills required to safely deploy and defend GenAI tools are not yet widespread in the public sector.
- Strategic alignment is urgent, with policy moving faster than skills. Without parallel upskilling efforts, federal teams risk deploying AI systems they don't fully understand or secure.
- CTEM is mission-critical and must become standard practice in the public sector, both to validate readiness and to keep pace with adversarial innovation.

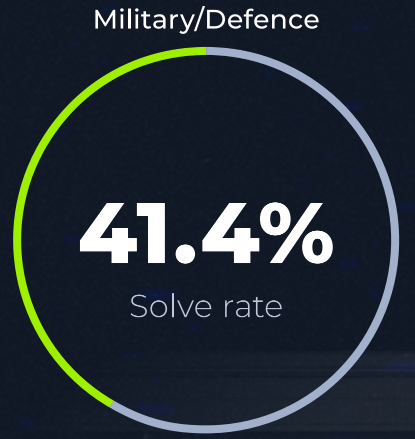

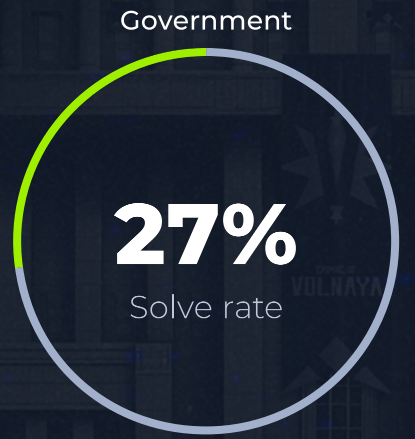

Finance

With deep regulation, constant exposure to fraud, and some of the largest digital footprints, finance should be leading the pack in cybersecurity. And in some ways, it is; Finance teams performed solidly overall (37.6%), showing strong capabilities in **Forensics (54.6%), Coding (51.4%)**, and **OSINT (71%)**.

But beneath that surface lies a concerning gap between perceived maturity and real-world readiness.

Solve rates in **Blockchain (10.1%), Web (19.2%), Secure Coding (20.8%), and Pwn (3.9%)** reveal that many financial institutions lack the hands-on skills needed to defend the systems they increasingly rely on—like online banking platforms, DeFi protocols, and custom payment flows.

Despite widespread investment in Web3 and embedded finance, **blockchain solve rates lagged behind sectors like healthcare and retail.**

This mirrors recent incidents, like the 2024 Euler Finance exploit, where poorly secured smart contracts led to major financial loss.

Low performance in Pwn could also point to a broader issue: minimal exposure to deep technical threats like ATM malware, memory corruption, or rootkits—risks that adversaries are already exploiting in the wild.

So what's driving the gap? One likely cause is strategic misalignment. Outdated, theory-heavy training, lack of incentives or time for continuous upskilling, and difficulty attracting elite cyber talent may all be weakening operational resilience in a high-stakes environment.

Ultimately, in a sector where uptime, trust, and transaction integrity are everything, complacency is pure risk.

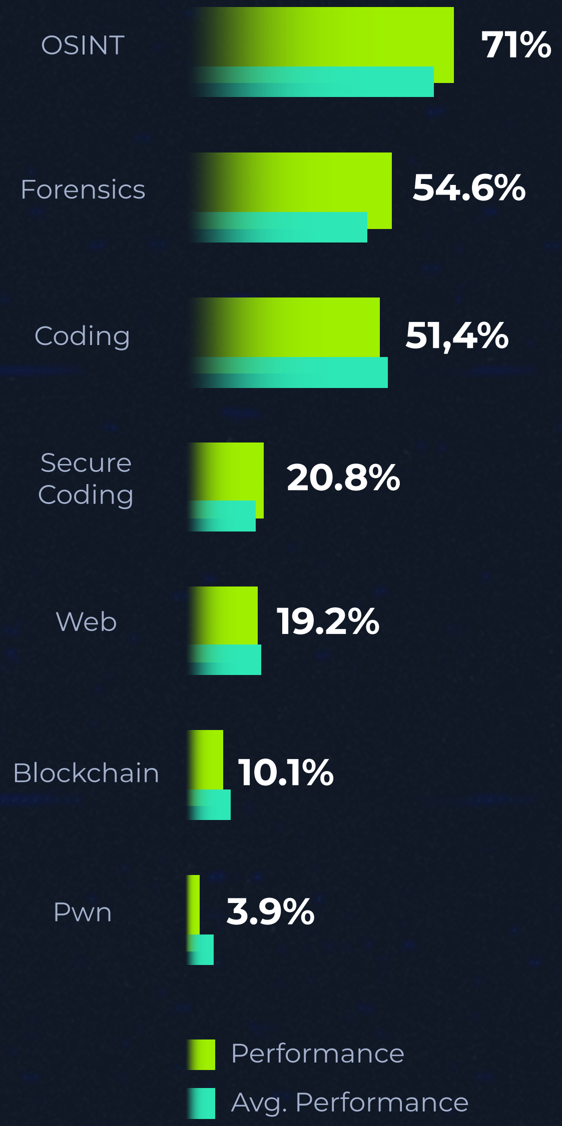

Finance

Key takeaways

- Detection is solid, but prevention is lagging. Strong Forensics and OSINT results contrast sharply with poor Secure Coding and Web performance.
- Blockchain fluency is critically low. Finance is investing in Web3, but internal security expertise isn't keeping pace.
- Offensive skills are underdeveloped. Pwn and Fullpwn results highlight a lack of practical exploit knowledge—vital for defending high-value systems.
- Strategic action is needed. Organizations should conduct business-aligned skills gap assessments, foster CTF-driven learning cultures, and invest in hands-on training tied to evolving tech roadmaps like AI, blockchain, and post-quantum crypto.

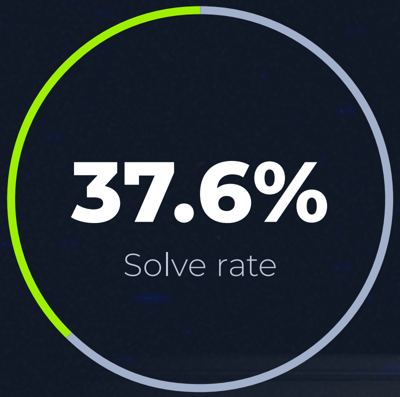

Healthcare

Healthcare remains one of the most frequently targeted sectors—but its cyber readiness in terms of skills and capabilities of technical teams in this space continues to lag. With an average solve rate of 27.5%, healthcare teams performed below average across most categories.

Web (15.6%), **Secure Coding (16.7%)**, and **Cloud** (**18.8%)** emerged as notably weak areas, reflecting persistent gaps in both application and infrastructure security. While **OSINT (76.6%)** and **AI (33.8%)** scores showed promise, foundational weaknesses in these core areas persist.

For a sector that handles life-critical systems and sensitive data, these gaps are more than technical—they're existential.

Key takeaways

- Healthcare teams must urgently improve secure development practices and cloud security hygiene to meet modern threat levels.
- Detection is improving, with high OSINT scores indicating growing maturity when it comes to reconnaissance and threat intelligence.
- There's an urgent need for SDLC and cloud hygiene. Addressing foundational gaps in development and infrastructure is crucial for managing modern threats.

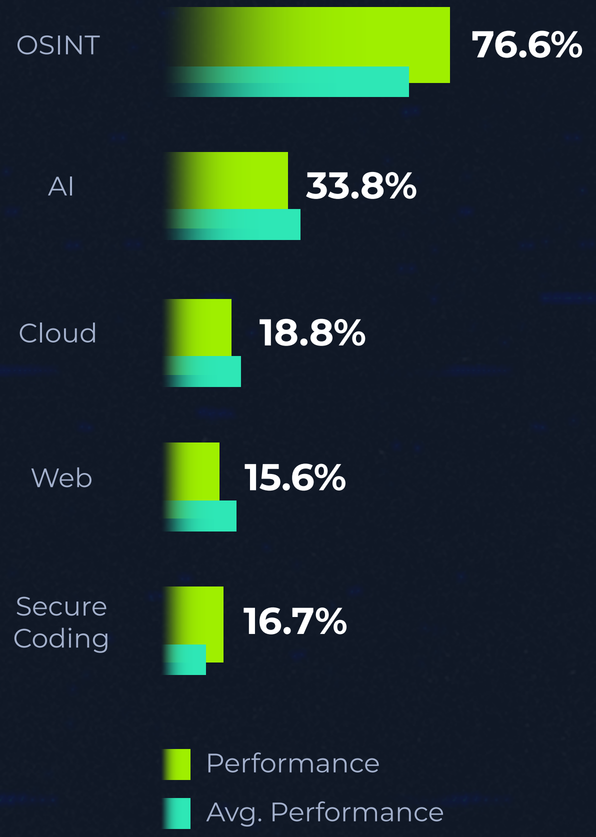

Education

Last year, Education was a chart-topping powerhouse. This year, it fell to the bottom, with an average solve rate of just 20.4%.

The biggest drop? **Secure Coding (0%)**—a stark result for a sector often seen as a talent pipeline for the industry. Other weak spots included **Cloud (1.3%)**, **Crypto (1.3%)**, and **Web (7.8%)**.

Bright spots like **Coding (60%)** and **OSINT (60.9%)** offer some reassurance, but overall, the data points to a widening skills gap, especially in real-world, adversarial challenge areas.

Key takeaways

- The education sector must reinforce secure coding fundamentals and modernize curricula to reflect today's cyber threats, not yesterday's theory.
- A 0% solve rate in Secure Coding is a major cause for concern, especially for a sector that feeds talent into the broader ecosystem.
- The gap between theory and practice is widening, meaning that curricula must evolve with real-world threat simulations embedded into training.

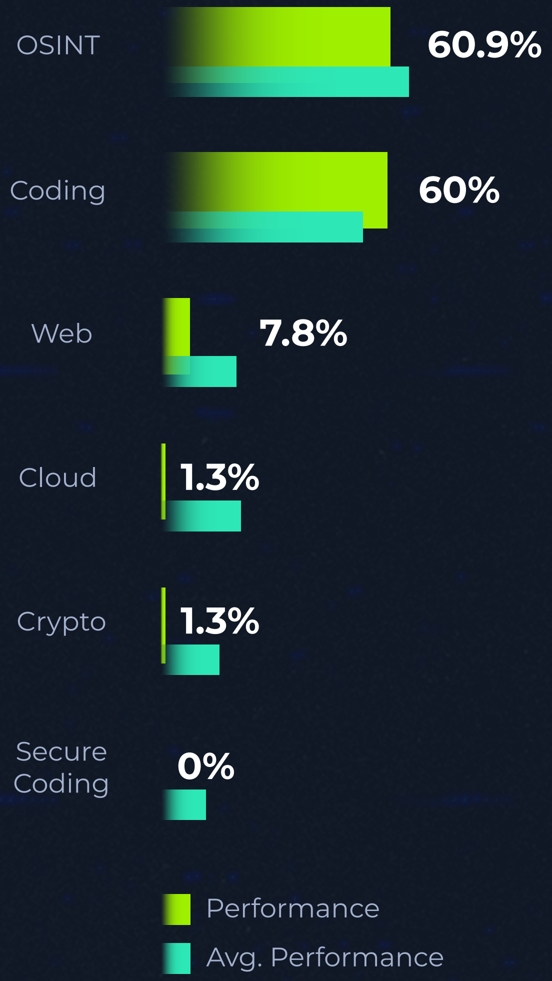

Manufacturing

Manufacturing teams delivered a quietly impressive performance this year, with a 35.6% overall solve rate—just behind Technology and Finance.

They led the way in **Reversing (29.1%)**, **Pwn (10.9%),** and **Secure Coding (30.3%)**, signaling strong capabilities in software exploitation and defense.

While **Web (18.2%)** and **Fullpwn (12.5%)** remained soft spots, the sector showed clear gains in embedded systems, ICS, and hardware-related categories—core areas of operational tech (OT) defense.

Key takeaways

- As manufacturing digitalizes, its cyber teams are evolving too—but offensive readiness and application-layer security still need reinforcement.
- Gaps remain at the application level. Defensive capabilities need reinforcement in web and full compromise scenarios, i.e. areas central to modern industrial IoT security.
- Manufacturing excelled in binary exploitation and reverse engineering, signaling readiness in embedded systems and OT security.

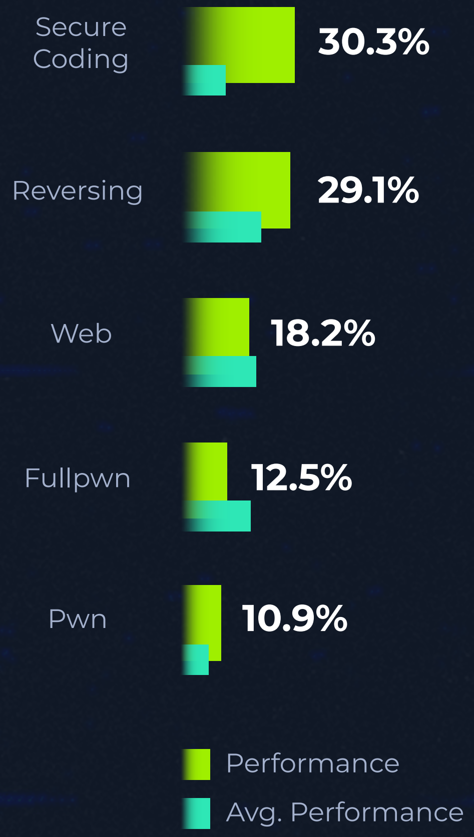

Energy and Utilities

Teams in this sector excelled in infrastructure-centric challenges, with strong results in **Forensics (55.6%)**, **Coding (42.7%)**, **OSINT (48.3%)**, and **ICS (21.7%)**. These scores reflect a solid foundation in legacy and operational technology (OT) environments—critical for defending industrial systems and critical infrastructure.

But the sector's performance fell sharply in modern domains. Solve rates were low in **Cloud (8%)**, **Web (6.7%)**, **Blockchain (2.2%)**, and **Secure Coding (4.4%)**, highlighting weak coverage in application-layer security. Even emerging areas like **AI (28%)** and **ML (20%)** show that while awareness is growing, practical skills remain limited.

With an overall solve rate of 31%, Energy and Utilities ranks mid-tier—but its polarized skill profile reveals a sector caught between legacy strengths and modern vulnerabilities.

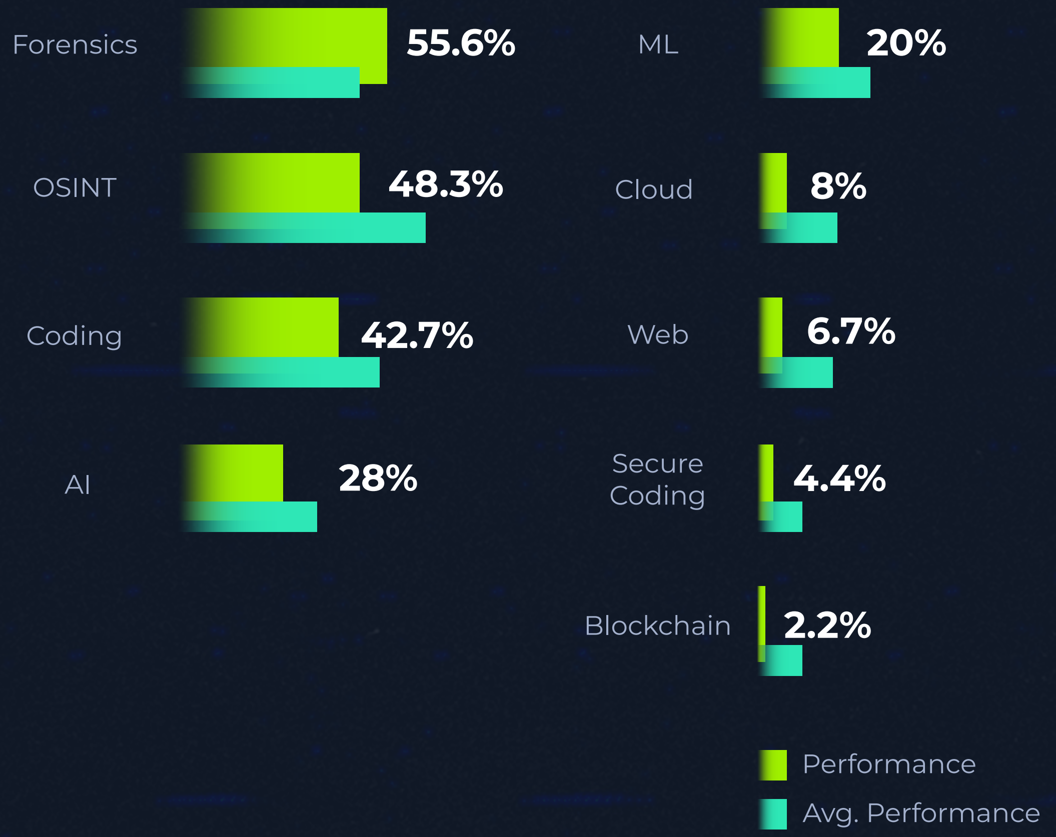

Energy and Utilities

Key takeaways

- Security maturity is high, but specialized. Forensics (55.6%) and ICS (21.7%) are strengths—but Cloud (8%), Web (6.7%), and Blockchain (2.2%) expose critical digital gaps.
- Modernization is a pressure point. As the sector digitizes, attackers will likely exploit weak points in cloud, apps, and development practices.
- The workforce needs a mindset shift. Secure Coding (4.4%) and Fullpwn (8.3%) performance show a need to move beyond perimeter defenses toward more adaptive, proactive security strategies.

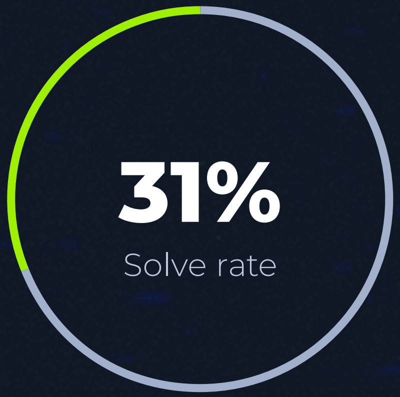

IT Services

IT Services teams delivered balanced, above-average performance across most categories, with notable strength in **Coding (53.3%)**, **OSINT (64.5%)**, **AI (38.3%)**, and **Cloud (37.6%)**. These results point to a broad, capable workforce well-versed in real-world cybersecurity demands.

That said, gaps remain in more specialized areas: **ICS (28.2%)**, **Hardware (29.7%)**, and **Pwn** (**9.8%**) underperformed compared to generalist categories. As more clients require OT and embedded systems support, these blind spots could limit future capability.

With a solid 35.3% average across all challenges, IT Services remains one of the most consistently high-performing sectors in this year's benchmark.

Key takeaways

- Generalists dominate, but niche gaps remain. While Coding (53.3%) and OSINT (64.5%) show technical maturity, Pwn (9.8%) and ICS (28.2%) reveal gaps in offensive and OT-centric skills.
- Client-driven demand must shape strategy. As industries digitize physical environments, IT services must expand fluency in hardware and infrastructure-specific threats.
- Training needs to go deeper, not broader. Strong generalist skills now require targeted investment in adversary emulation, OT defense, and low-level exploitation to stay ahead.

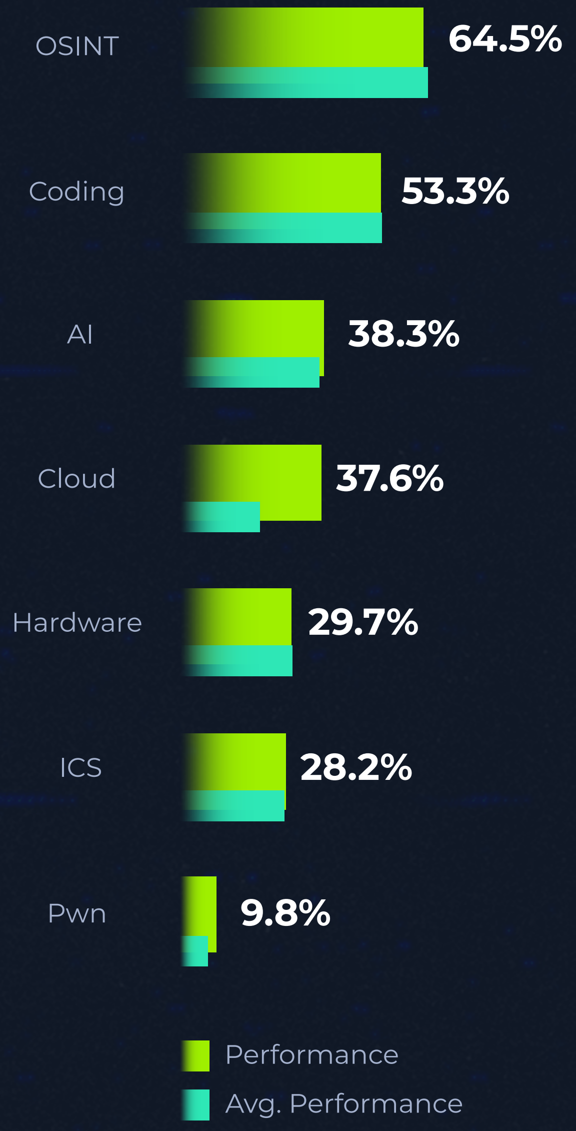

From benchmark to battle-ready

Turn insights into risk reduction

CTFs like The Global Cyber Skills Benchmark surface the weak spots in your skill set. Meanwhile, CTEM uses controlled threat emulation to measure real-world impact. Think of it as the bridge between Turn insights into risk reduction knowing a gap exists and understanding what it costs.

Say your team underperforms in Web and Fullpwn. Benchmarking highlights that weakness. CTEM simulates the breach: How far could an attacker go? How long would it take to detect? What would be lost?

Benchmarking reveals the gaps. CTEM quantifies the risk.

Here's the path a high-performing team would take:

- **Benchmark** with real-world performance data.
- **Emulate** threats targeting known weaknesses.
- **Quantify** exposure—detection time, escalation paths, data at risk.
- **Remediate** with targeted training, then re-test.

Working together, benchmarking and CTEM create a feedback loop: insight → emulation → action → repeat. In a landscape of constant exposure, that loop is where you can get your competitive edge.

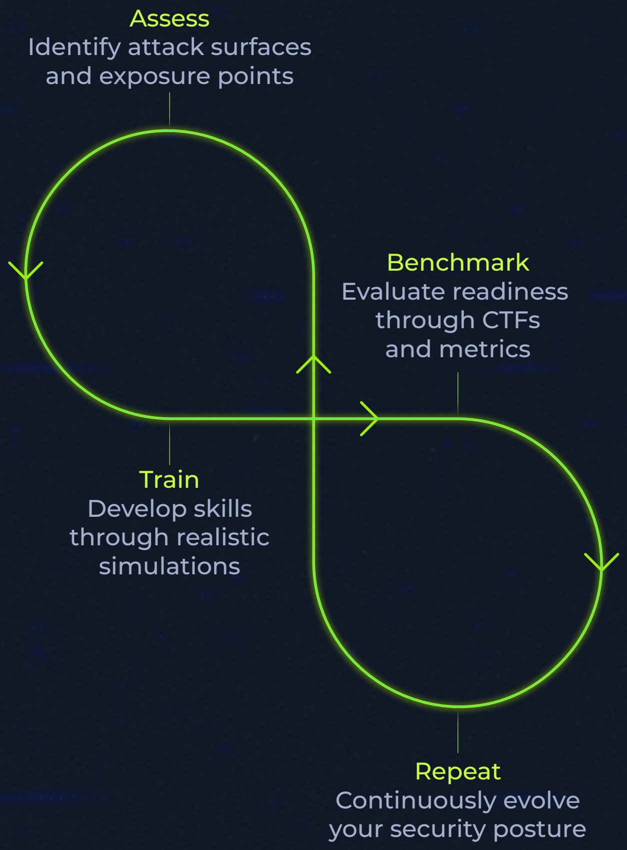
Assess
Identify attack surfaces and exposure points
Benchmark
Evaluate readiness through CTFs and metrics
Train
Develop skills through realistic simulations
Repeat
Continuously evolve your security posture

Conclusion and next steps

Cyber readiness isn't a checkbox.
It's a performance discipline.
The Global Cyber Skills Benchmark doesn't just tell you where you stand—it shows where you're vulnerable, where your industry is headed, and how to get ahead of evolving threats.

If you're in finance, healthcare, government, or education, the gaps are no longer theoretical. They're measurable. And they're exploitable.

What's next? Act on what you've seen. Map your skills. Set your benchmarks. Build your training around reality, not guesswork. Because attackers aren't waiting, and your readiness is the only edge that matters.

Kick things off by:

- → Reviewing your industry's weakest domains (with special attention on Secure Coding, Web, and Cloud)
- → Mapping your team's skill levels against your peers
- → Using Hack The Box to assess, train, and track real progress

What teams say about our global CTF

“

It was a fantastic and well-rounded experience. One of the highlights for me was getting to apply memory forensics—something I had only recently learned while preparing for a session I delivered.

Using those skills in a live challenge really **reinforced the learning** in a practical way. The CTF struck a great balance between technical depth and hands-on application, which made it both fun and rewarding.

I work in network security, so I particularly appreciated the challenge involving PCAP analysis—it aligned closely with the kind of traffic investigation I do in my day-to-day work. Beyond that, I spent a lot of time **collaborating with my team**, discussing our approaches to different challenges. Even when I wasn’t directly solving the challenge, I learned a lot just by following their thought process and strategy; it was a great team learning experience.

Ravi Shankar Krishna
Security Technical Consulting Engineer at CISCO

”

“

I’m proud to share that I competed for Microsoft’s team in the Hack The Box Company CTF (Global Cyber Skills Benchmark CTF 2025: Operation Blackout) and we finished 8th out of 800 teams! Thank you to all the amazing Microsoft employees who competed—your energy, creativity, and collaboration made this possible.

We’re already gearing up for next year and ready to claim 1st place!

Jovan Pavlovic
Software Engineer at Microsoft

”

The Global Cyber Skills Benchmark 2025

HACKTHEBOX