



Research Brief

The AI-Accelerated Cyber Team

What three years of competition data reveal about speed, coverage, and the changing value of expertise.



Three years of cyber performance under pressure

The Global Cyber Skills Benchmark is the world's largest hands-on cyber competition for business security teams, built to test practical capability under pressure against a global field.

Across 2024, 2025, and 2026, teams faced realistic cybersecurity challenges spanning a broad range of difficulty. Looking across the three events gives us a view of how competition performance changed over time.

Alongside that series, Hack The Box also ran the [AI-Augmented vs Human-Only Cybersecurity Performance Benchmark](#), based on the NeuroGrid CTF. That research adds a separate view of how AI augmentation changes performance when humans and AI work together.

Taken together, the four events provide context for the changes explored in this report, from broader capability across teams to faster execution and the growing role of AI in cyber work.

2024

HTB Business CTF

The Vault of Hope

2025

Global Cyber Skills Benchmark

Operation Blackout

2025

NeuroGrid CTF

AI-Augmented vs Human-Only Cybersecurity Performance Benchmark

2026

Global Cyber Skills Benchmark

Project Nightfall

More teams are building the capability to clear the entire CTF



What competition data from 2024 to 2026 suggests about breadth and operational coverage.

A high score can hide gaps. A team may dominate several categories and still stop when the work moves into an uncovered domain. Full completion asks whether the team can cover everything the competition presents.


Two teams cleared the board in 2024. Three did so in 2025. By 2026, 15 teams reached full completion.

Board size and challenge mix changed across the period, so full completion is not a fixed measure of technical difficulty. It does show that more teams developed the breadth, coordination, and endurance needed to keep progressing across the full event.

No workforce needs every practitioner to master every domain. It does need enough team coverage to prevent one unfamiliar system, absent specialist, or capability gap from stopping the response.

Full-board completions

Completion of all challenges in the CTF.



Board size and challenge mix changed across years. Counts show full completion, not relative difficulty.



Productive execution moved forward

Teams reached valid outcomes more than 12 hours earlier by 2026.

The midpoint of recorded solve activity moved forward across the 3-year period.

In 2024, the median recorded solve arrived after 1,564 minutes. By 2026, it arrived after 827 minutes. That is a shift of 737 minutes, or just over 12 hours.

This does not mean every practitioner became 47% more skilled. The challenge mix, board design, competitor population, preparation, and available tooling all changed across the period.

It does show that useful execution happened earlier.

The cross-year data cannot isolate a single cause. Better preparation, reusable tooling, platform familiarity, and deeper specialist knowledge all played a likely part.


AI should be treated as a leading accelerator within that mix. Separate HTB benchmark research found that AI-augmented teams worked faster than human-only teams, with the speed advantage increasing among the strongest operators.

Read together, the evidence points to a clear shift. AI is compressing the work around expertise. It can shorten research, coding, troubleshooting, and the path to a first viable answer. The practitioner still decides where to look, what to trust, and whether the result is valid.

The next question is not whether AI replaced expertise. It is how much faster expertise became when AI entered the workflow.

Median recorded time-to-solve

Minutes from event start. That is a shift of 737 minutes, or just over 12 hours, across a comparable mix of challenge difficulty.



**737 MINUTES
EARLIER**

**2024 TO 2026
MIDPOINT SHIFT**

Board size and challenge mix changed across the period. These figures are directional performance signals, not a fixed comparison of individual skill.



AI amplified prepared teams. It did not replace them.

Project Nightfall and the HTB benchmark research point to the same human-led operating model.

The Global Cyber Skills Benchmark 2026: Project Nightfall and [the Hack The Box AI-Augmented vs Human-Only Cybersecurity Performance Benchmark](#) point to the same answer. AI accelerated prepared teams while judgment and ownership remained human.

Ninety-three agent accounts appeared across 54 teams. They represented 2.7% of registered accounts, 4.2% of submitted flags, and 4.6% of awarded points. Agent accounts were more common near the top of the leaderboard, appearing in 17 of the Top 25 teams.

That association does not prove that agents caused stronger performance. It does show that leading teams were already incorporating AI into the way they worked.

The Hack The Box AI-Augmented vs Human-Only Cybersecurity Performance Benchmark, based on the NeuroGrid CTF in 2025, tested the same operating question more directly. It compared AI-augmented teams with human-only teams on the same challenge set.

The findings converge. Across all active teams in the HTB benchmark research, AI-augmented teams recorded a 3.2 times solve-rate advantage. That advantage narrowed to 1.69 times among the Top 5%, while the speed advantage grew. Elite AI-augmented teams completed challenges three to four times faster, yet the strongest human team completed all 36 challenges compared with 32 for the strongest AI team.

AI therefore creates the greatest value when it is directed by people who already understand the problem. It can compress research, coding, troubleshooting, and first-pass analysis. The operator still has to choose the path, challenge weak output, and decide whether the result is valid.

AI is not replacing cyber expertise. It is increasing the speed at which strong expertise can be applied.

PROJECT NIGHTFALL

Designated agent participation

93	DESIGNATED AGENT ACCOUNTS	54	WITH AN AGENT ACCOUNT
46	ACTIVE AGENT ACCOUNTS	2.7%	SHARE OF REGISTERED ACCOUNTS
4.2%	SHARE OF SUBMITTED FLAGS	4.6%	SHARE OF AWARDED POINTS
17 OF 25	TOP 25 TEAMS WITH AN AGENT ACCOUNT	33 OF 100	TOP 100 TEAMS WITH AN AGENT ACCOUNT

HTB BENCHMARK RESEARCH CONFIRMS THE PATTERN

3.2x	1.69X	3 TO 4X FASTER
AI solve-rate advantage across all active teams	Advantage among the Top 5%	Elite AI-augmented teams

Project Nightfall measures agent presence and credited activity. The HTB AI-Augmented vs Human-Only Benchmark directly compares AI-augmented and human-only performance. Neither benchmark supports removing human oversight.

Sources: Project Nightfall 2026 event data. Supporting evidence from the Hack The Box AI-Augmented vs Human-Only Cybersecurity Performance Benchmark Report, based on the NeuroGrid CTF, November 2025.



Two benchmarks. One workforce mandate.

AI changes the value of human expertise differently across the workforce.

Project Nightfall and the Hack The Box AI-Augmented vs Human-Only Cybersecurity Performance Benchmark examined different parts of the same shift. Project Nightfall measured designated agent participation and credited activity in a large, live competition. The HTB benchmark research directly compared AI-augmented teams with human-only

teams on a shared challenge set. Both point to the same workforce model. AI can raise output across the field, but its value changes with operator experience. The capability advantage narrows as human expertise rises. The speed advantage becomes more pronounced when skilled practitioners use AI to accelerate work they already understand.

AI SPEED ADVANTAGE GROWS WHEN EXPERTISE IS ALREADY STRONG

EARLY CAREER

**Build judgment,
not dependency**

AI can help less experienced practitioners complete work they might not reach alone. But output can rise faster than understanding. Keep hands-on work in the development path. Teach cyber fundamentals and AI direction together. Require practitioners to explain, test, and validate AI-assisted results.

- **Cyber fundamentals**
- **Hands-on investigation**
- **Mandatory validation**
- **AI-assisted learning**

MID CAREER

**Multiply capable
operators**

Mid-career practitioners can convert AI support into operational throughput because they have enough context to direct the tool and enough experience to challenge weak output. Use AI to compress research, scripting, triage, documentation, and repeatable analysis. Keep escalation, evidence quality, and final decisions with the practitioner.

- **AI-assisted triage**
- **Research and scripting**
- **Faster investigation**
- **Human-owned decisions**

ELITE

**Retain expertise,
then accelerate it**

At the elite tier, the capability gap narrowed while the speed advantage widened. This makes experienced practitioners more valuable, not less. Keep the more complex work human-led. Pair top practitioners with AI and protect the talent capable of recognising when the tool has reached its limit.

- **Parallel hypothesis testing**
- **Complex validation**
- **Human-led response**
- **Retention and succession**

Build judgment early. Multiply capable operators in the middle. Retain and accelerate elite expertise at the top.



Cyber Mastery: Community Inspired. Enterprise Trusted.

Hack The Box is the leading cyber readiness platform for the agentic era, battle-testing and upskilling both humans and AI agents to enhance organizational cyber resilience. Trusted by the Fortune 500, government agencies, and Managed Security Service Providers (MSSPs), the platform delivers threat-informed learning paths consisting of real-world scenarios in gamified labs and live-fire simulations that build and validate offensive and defensive cyber capabilities. With a loyal community of more than 4.3 million members and more than 800 enterprise customers, Hack The Box empowers teams and intelligent systems alike to strengthen cyber defenses and reduce breach risk effectively.

hackthebox.com

