The Death of the CTF: How Agentic AI Is Reshaping Competitive Hacking

Jacob Krell | February 2026

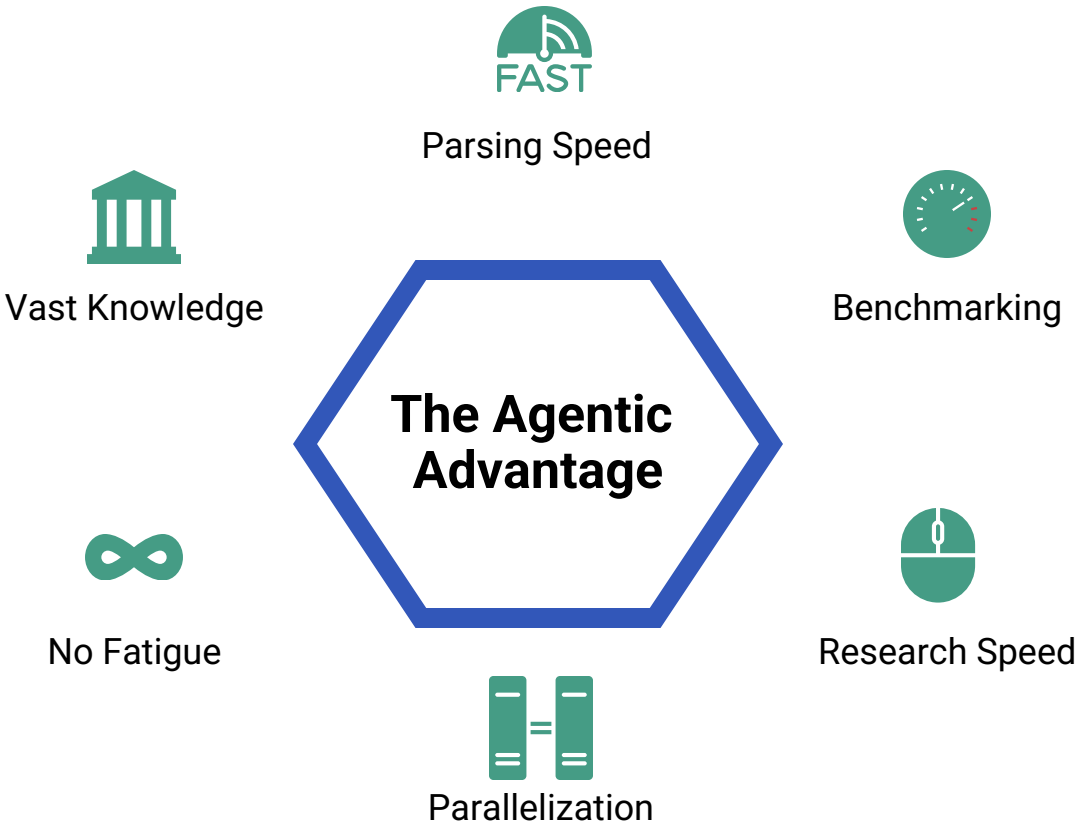

Agentic AI systems are compressing competitive hacking timelines faster than the cybersecurity community has acknowledged. This shift redefines the human element in CTFs.

The competitors who will succeed in an AI-augmented CTF landscape are not those who resist the technology but those who learn to leverage it. Directing AI agents toward the tasks where they hold categorical advantages while focusing human effort on the areas where intuition, creativity, and adversarial reasoning remains superior. The skill being tested is no longer purely operational. It becomes architectural, and the question is: how effectively can a competitor design, configure, and orchestrate an AI system?

The Three Waves of AI Offensive Capability

Wave 1: LLM Assistance Early research showed LLMs could assist with security tasks. Palisade Research achieved 95% on the InterCode-CTF benchmark, proving capabilities were underelicited.

Wave 2: Real-World Exploitation LLM agents began autonomously exploiting real CVEs. Google's Big Sleep discovered the first publicly documented AI-found exploitable bug in real-world software.

Wave 3: Competitive Dominance In 2025, the Cybersecurity AI agent won the Neurogrid CTF and reached Rank 1 at the Dragos OT CTF, outperforming thousands of human teams.

Compressed Timelines Analysis of 423 HTB machines shows root blood times declining 16% per year. The hardest machines saw a 67% reduction in solve times in the Post-LLM era.

Privilege Escalation Acceleration The privilege escalation phase is compressing faster than the foothold phase, dropping by 41% at Hard difficulty as agents handle post-foothold enumeration perfectly.

Platform as Benchmark Public CTF platforms now serve as de facto benchmarks for national AI cyber capabilities, providing real-time signals about global AI ecosystems.

Longitudinal Analysis of HTB First Blood Data

Analysis of 423 Hack The Box machines released between 2017 and 2025 reveals significant compression in solve times:

16%
Annual decline in blood times across all difficulties

67%
Reduction in median root blood for Insane difficulty

27%
Reduction in median root blood for Hard difficulty

47%
Reduction in median root blood for Easy difficulty

The greatest reduction observed was the time from user blood to root blood, or the privilege escalation step.

41%
Compression of privilege escalation time at Hard difficulty

A sustained 16% annual reduction implies a sixfold decrease in time-to-compromise within a decade, breaking existing security models.

The End of Unaided Cognition

The CTF as a purely human competition between operators is ending. What replaces it is a richer, more complex competition between the humans who design AI systems and the AI systems they build. The organizations, competitors, and governments that recognize this transition earliest and adapt most effectively will define the next era of cybersecurity. The ones that pretend nothing has changed will be left wondering why the scoreboard no longer makes sense.

Read the Full Synopsis Here: https://suzulabs.com/suzu-labs-blog/the-death-of-the-ctf-how-agentic-ai-is-reshaping-competitive-hacking