Can AI Do Novel Research?

Meet the HTTP Terminator

Executive summary

Everyone accepts that AI can find vulnerabilities, but many claim it can't do original research. To challenge that assumption, I built the HTTP Terminator: an autonomous system that invents new attack techniques and uses them to hack live websites at scale. I aimed it at the topic I'm most qualified to judge: HTTP Desync Attacks (AKA HTTP Request Smuggling). I've spent four years and four Black Hat/DEF CON talks researching these attacks, so if it tried to falsely claim it had discovered a new technique that was, in fact, already public knowledge, I'd know.

It worked. The system follows my own research process: ideation, evaluation, weaponization, and cascade. It read 138 technical specifications, broke them into 15,000 fragments of inspiration, and generated 30,000 unique attack vectors. It then tested these continuously against live targets authorized through bug bounty programs, confirming roughly 700 vulnerable systems and autonomously proving real-world impact. Compromised targets included government infrastructure and multiple banks (including one from which it accidentally stole a live API key). The breaches were traced back to flaws in widely deployed enterprise products including BeyondTrust, Apache Traffic Server, and Citrix NetScaler.

The deeper finding concerns the human role. Fully autonomous research is real - I proved you can build the loop and watch the findings rain - but the system's true power is unlocked by putting a researcher back in at exactly one point: the discovery cascade, where each finding becomes the seed for the next. This yielded multiple significant discoveries, including a genuinely new class of vulnerability I call Shared-Parser Confusion.

This inverts the accepted narrative by showing an expert can be a massive amplifier for an AI research system. I'll open-source both the HTTP Terminator and a blueprint for researchers to turn their own methodology and instincts into an autonomous research weapon.

You can find the full whitepaper at
https://portswigger.net/research/can-ai-do-novel-security-research

This research was presented at Black Hat USA 2026 and DEF CON 34, and the linked page will be updated with the presentation recording once it's available - follow PortSwigger Research on X, LinkedIn or RSS to get notified when it lands.

Author biography

James 'albinowax' Kettle is the Director of Research at PortSwigger, the makers of Burp Suite. He's best known for pioneering influential web attack techniques like HTTP Desync Attacks - publishing novel research at Black Hat USA for ten consecutive years. These discoveries are underpinned by an automation-first mindset, inventing brutally effective tools that yield new attack classes and often become industry standard.

Popular attack techniques that can be traced back to his research include HTTP Request Smuggling, web cache poisoning, the single-packet attack, server-side template injection, and password reset poisoning. His extensive tool legacy includes introducing OAST via Burp Collaborator, bulk parameter discovery via Param Miner, and billion-request attacks with Turbo Intruder. He's also the designer behind many of the topics and labs that make up the Web Security Academy, runs the annual Top Ten New Web Hacking Techniques project, and serves on the Black Hat Europe review board.