The Next Advancement with OpenAI: Zscaler Steps Up To GPT-5.5-Cyber

May 11, 2026

Cybersecurity has always been a race. Attackers look for a way in, defenders try to stop it, and whoever moves faster wins. For thirty years, both sides moved at human speed. With AI, the game has changed, making AI security an urgent priority. The most advanced AI can now find weaknesses in software and turn them into real attacks in minutes. It does not sleep, get tired, or miss details. The advantage cybersecurity defenders used to have, knowing their own systems better than anyone, is gone.

OpenAI has been thoughtful about what to do with that capability. Rather than release their most powerful cyber-focused AI to anyone, they created the Trusted Access for Cyber program, or TAC, to make those capabilities more useful for verified defenders. Zscaler has been a proud member of the TAC program since the beginning, starting with GPT-5.4-Cyber. Today, we are stepping up to GPT-5.5-Cyber, the next advancement from OpenAI.

If your security strategy still depends on patching faster than the adversary can find the bug, you have already lost. The cycle time has collapsed. The winning move is to stop applications from being visible to the adversary in the first place, and to bring an AI of equal caliber to the work of finding what humans miss. That is what TAC delivers, and that is what Zscaler was built for.

Two Things Defenders Need

There are two things a modern defender needs, and they have to work together.

1. An architecture that removes the target - If an application is invisible to the internet, it does not matter how clever the model is on the offensive side. There is nothing to scan, nothing to reach, and nothing to weaponize. This is the entire point of Zero Trust, and it is the principle the Zscaler Zero Trust Exchange has operationalized for nearly two decades. Users do not connect to the corporate networks. Applications do not have public IP addresses. Every session, human or machine, is securely brokered one-to-one against a verified identity. We have a saying inside Zscaler: if you are reachable, you are breachable. You need to make yourself unreachable.

2. Frontier AI Working on the Defender's Side - Removing the target is necessary, but it is not sufficient. There are still bugs in code, still misconfigurations in the cloud, still phishing in inboxes, still anomalies in the 500 billion daily transactions we see on our platform at Zscaler. Sorting through that volume and acting on it at machine speed is what a model like GPT-5.5-Cyber is built for. Make the application, data and workloads invisible to attackers via the Zero Trust Exchange platform. Then leverage a frontier model with what remains. That is the combination.

Advancements Leveraging GPT-5.5-Cyber

We have spent a significant amount of time inside TAC with GPT-5.4-Cyber, and we have learned where a model of this caliber pays off most. GPT-5.5-Cyber sharpens the work in three places.

- **Hardening the AI our customers are racing to deploy:** AI tools can be tricked or manipulated into giving up sensitive data. Our AI Red Teaming capabilities (formerly SPLX) has used OpenAI since early 2024 to generate the attack sequences that test and harden customer AI across text, voice, and images. With GPT-5.5-Cyber, when Zscaler finds a weakness, it generates an optimized fix in the same loop, a working first step toward a stronger security posture. The same engine scans the code behind complex AI agents and the tools they connect to, and it powers our Agentic Radar open source project, which fueled the largest OpenAI hackathon last year in Warsaw.
- **Building safer software inside Zscaler:** GPT-5.5-Cyber works alongside our engineers across the way we design, build, and ship our solution, validating security thinking at the design stage, assisting with code reviews, finding vulnerabilities at machine speed, and probing finished software the way an attacker would. Problems get caught and fixed long before any customer is exposed to them.
- **Faster investigations without losing the human in the loop:** Inside Red Canary, our managed security service, OpenAI-powered agents do the tedious context-gathering that traditionally overwhelms security teams, while our human experts make the final call. The same foundation supports continuous red teaming, prompt hardening, and AI asset analysis across our broader platform — faster, more accurate defense at the speed our customers need.

We are also committed to giving back for the benefit of everyone. As part of TAC, the lessons we draw from running GPT-5.5-Cyber at this scale flow back to OpenAI and the broader coalition. Defenders win as a community, or not at all.

A Pattern the Industry Keeps Forgetting

When the cloud arrived, most of the industry believed existing defenses would still work. They did not. When everyone moved to phones and online apps, most believed traditional remote-access tools would adapt. They did not. Each time, the winners were the companies that accepted the world had changed and built for what was coming.

Zscaler was one of those companies. More than 17 years ago, we made a bet that the future would not be defended by firewalls and VPNs guarding a network-based perimeter. We built our cybersecurity platform on Zero Trust instead, an architecture where applications are invisible to the internet, and every user or device is securely verified for one-to-one connections. That bet was right for the cloud era, and it is also right for what's next.

AI is the next of these moments, and it is moving faster than the ones before it. Attackers already have the technology. Through our partnership with

OpenAI, defenders do too. But the architecture matters even more now than it did before. A frontier-class AI pointed at an exposed application is a problem no patch cycle can keep up with. A frontier-class AI pointed at an application it cannot even see is a different conversation entirely. That is the conversation Zscaler has been preparing the industry to have since 2008.

The question for every leader is simple. Keep guarding a front door the attackers no longer use, or move to a security model built for what is coming. The longer the decision waits, the more likely a breach makes it for you.