Inside the Network That Hackers Try Hardest to Break:
A Look at Cybersecurity's AI-Driven Future

Most professional conferences hand their guests the hotel Wi-Fi password and call it a day. Black Hat doesn't. Before each of its annual gatherings in Las Vegas, London, and Singapore, the organizers build a network from scratch — switches, access points, firewalls, monitoring sensors — and stand it up inside the venue. The Network Operations Center, or NOC, then has to defend that network, in real time, against thousands of the world's best hackers. Some are external adversaries trying their luck. Many are attendees the conference has explicitly told to attack the infrastructure as part of the curriculum.

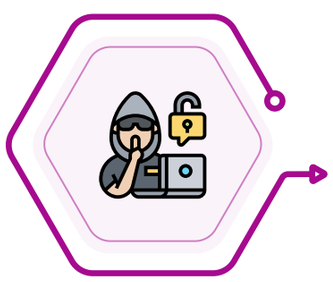

This year's Singapore edition, held April 21st to 24th, arrived against an uncomfortable backdrop. Major technology companies have been announcing that their AI models can now outperform almost every human hacker on Earth. Anthropic's Mythos, the most prominent of these, has reportedly identified severe vulnerabilities across every major operating system and web browser. For ordinary users, this feels like a sudden lurch into a new era. For the people running Black Hat's NOC, it's the moment they've been preparing for.

Defending Black Hat Is Not Like Defending Anything Else

Neil "Grifter" Wyler has run the NOC for twenty-four years. For most of that time, he's worked alongside his colleague Bart Stump. The job, Wyler says, is orders of magnitude harder than ordinary corporate cybersecurity — so much so that when the head of cybersecurity for the Paris Olympics needed a working model for his own security operations center, he spent a week embedded with the NOC at Black Hat London.

Part of the difficulty is scale. A typical company sees one or two active attackers at a time. Black Hat sees thousands, and many of them are working through exploits that world-class instructors taught them the day before. The other part is filtering. The NOC has to let legitimate coursework happen — that's the whole point of the conference — while distinguishing it from genuine malicious activity. The line between the two is sometimes invisible.

What they see ranges from the absurd to the alarming. One attendee was running a weather app that quietly leaked their GPS coordinates. Another was using a pet-feeding app that anyone with basic skills could have hijacked. Eighty-one unique adult-website domains made the visit log.

The Real Attacks Hidden in the Noise

The same tools that catch compromised pet feeders also catch the serious stuff. A few years ago, a participant used the conference network to hack an American water-treatment facility. Wyler and Stump won't say much more about that one. Another attacker hid inside the legitimate hacker traffic to go after government websites and payment systems. The NOC traced him, sent him a message reminding him that doing illegal things from Black Hat was still illegal, and watched as he closed his laptop and walked off.

Outside attackers try too. The moment the registration server comes online, attacks begin — including, recently, traffic that appeared to originate in Romania. "It would be a feather in their cap to take down Black Hat," Wyler says.

Where AI Has Already Taken Over

AI has been part of the NOC's defenses for years, deployed against automated bots and human attackers alike. What's changed is the bots themselves. "The attacks have gone from taking a week to a day to hours or minutes," Wyler says. To keep pace, the team has assembled a stack of AI tools to fight fire with fire.

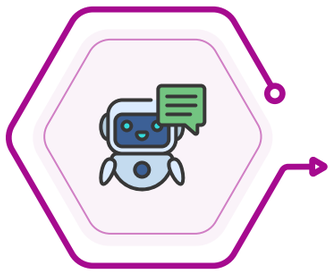

One of them, an internal chatbot the team calls Trevor, converts plain-English questions into the database queries the NOC's complex monitoring stack actually needs. That sounds modest. In practice, it gets new freelance analysts productive in hours instead of weeks. A second tool watches the patterns of encrypted beacons — the tiny, regular check-ins that compromised devices send to attackers' command servers — and uses machine learning to separate them from the millions of legitimate connections happening every day on the same network.

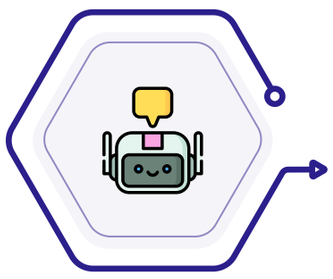

That second tool is how the team caught the compromise of a Taiwanese journalist's laptop at the conference. Among all the ordinary traffic, his machine was contacting an unfamiliar server at a metronomic cadence, returning at intervals no legitimate app would produce. A third tool — an AI agent — then profiled the device, cross-referenced clues from the network against publicly available information, and identified the owner. The team confirmed the match through the conference's registration database, wrote up a report, and informed both the journalist and his organization.

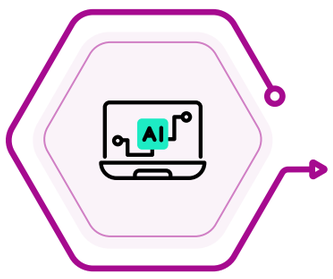

This isn't an isolated case. Stump says the NOC has seen the same pattern at multiple Black Hat conferences: Taiwanese attendees arriving with already-compromised devices, traffic from those devices flowing back to China. The implication is one no security professional needs spelled out. Nation-state and criminal AI-powered attacks are about to get worse.

The Next Two Years

The team thinks the AI race is just starting. For Wyler, the new wave of vulnerabilities surfaced by tools like Mythos — some of them latent for decades — is welcome, not alarming. "We now know they're there," he says. The flaws were always present. They simply weren't visible.

Stump is more cautious about the transition itself. The next two years, he expects, will be turbulent. More flaws will be found. More breaches will happen as companies feed sensitive data into AI systems without thinking carefully about what they're exposing. More insecure code will be written by people who trust AI assistants further than they should. If that transitional period is handled with discipline, a new equilibrium may eventually emerge that looks a lot like the one we're leaving behind — uneasy, imperfect, but workable.

One thing, Stump says, is certain. "In a year there will be a new AI model that makes Mythos look like a toddler with a keyboard."

The cybersecurity world has spent the past decade preparing for AI to change the game. At Black Hat, the change isn't theoretical anymore. It's the network they're already defending — and the new tools they're already using to defend it.