The "Good Guy Hackers" Helping Education Companies Test Their Cybersecurity Practices

The Process of "Red Teaming" Allows Vendors to Probe Their Vulnerabilities

By Alexandria Ng — August 30, 2024

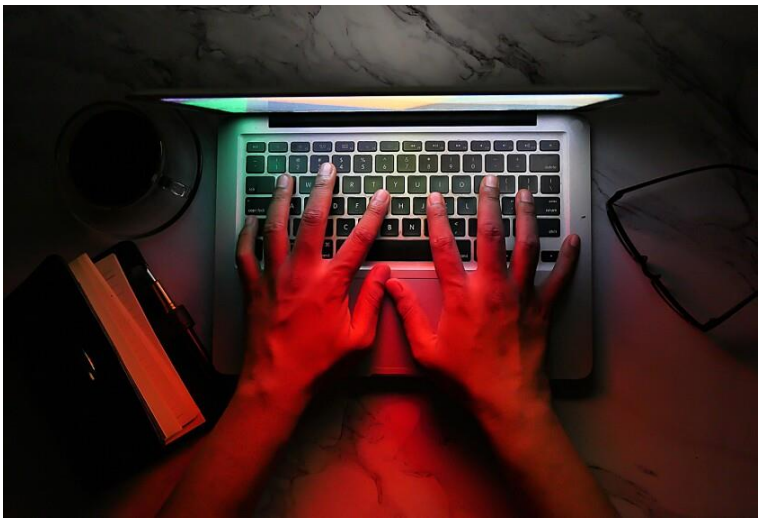

To fight against bad actors, sometimes you need to get into their mindset.

When it comes to cybersecurity, one way that education technology companies are fighting back against attacks is through a process called red teaming. That's when a group of security experts play the part of malicious actors to exploit weaknesses in a system and help organizations build up a stronger defense against real-life attacks.

Education companies' interest in red teaming comes as cybersecurity has emerged as a major concern for these providers, and for the school districts they serve, which have faced a growing array of cyberthreats in recent years.

Between 2016 and 2022, there were 1,619 cybersecurity-related incidents reported in U.S. K-12 public schools and districts, as tracked by the K12 Security Information eXchange, a national nonprofit dedicated to helping schools defend against emerging cybersecurity threats.

Red teaming is seen by education companies as a way to not only protect their organizations' own data, but also the information they may have responsibility for managing in school districts.

"There's been an increase in the number of attacks that are taking place, and it's having a real impact on operations and loss of data," said Will Sweeney, managing partner and

founder of Zaviant, which helps K-12 and higher education institutions build out their data security and privacy programs.

The education sector has historically "underinvested in this particular area," he added, but the need for stronger cybersecurity practices has risen with "increased scrutiny and regulatory oversight."

The number of education companies undergoing red team processes is still not very high. According to Cobalt Offensive Security Services, a provider of red team services, only 10 to 20 percent of their customer base comes from the ed-tech sector.

Those organizations represent only "a minority [of] our customers," said Caroline Wong, chief strategy officer for Cobalt. "I encourage [vendors] to research security attacks that have been conducted on their peers and on their competition and ask themselves what they would do in that situation if that type of attack happened to them."

EdWeek Market Brief spoke to officials in the cybersecurity space to discuss how red teaming works and the benefits it can provide in not just protecting internal and external-facing systems, but strengthening protections for districts and building trust between companies and school systems.

Process Breakdown

Red team exercises aim to simulate a cyberattack to assess a system's vulnerabilities and see if proper protections are in place to prevent those attacks from succeeding.

The actual team of "hackers" on a red team project will vary depending on the nature of the test. During the exercise, the security experts will use a variety of tactics to try to penetrate an entity's system.

The exercise typically begins with the hackers conducting reconnaissance. That could take the form of a black-box strategy, in which the red team comes in blind, with no knowledge of an organization's internal systems.

In a white-box strategy, members of a red team may be set up with login credentials to then go after a system's architecture and code. The data collected through either approach will be used later by the red team to launch an offensive attack.

The education organization being tested won't know when the attack is coming. It could happen within weeks or even months.

At the end of the test, the red team will provide a post-breach report and a briefing, in which the organization conducting the attack will explain to the company's internal teams what

vulnerabilities were found, and what next steps should be taken to fortify the company's defenses.

Recommendations for improvements could include steps such as training employees on how to avoid phishing attacks, how to fine-tune tools that detect and respond to cyberthreats, and how to shore up weak firmware.

It's important to find the right providers to perform this service, Zaviant's Sweeney said, as a poorly performed red team exercise could potentially have an impact on system operations and degradation of functionality.

"You want someone who is using a well-defined methodology because there's the potential for systems to be brought down to a point where that system is unusable because of the attack," he said.

Red Teaming at Work

This summer, K-12 software company PowerSchool enlisted a third-party red team service provider with the goal of fortifying PowerBuddy, its AI assistant designed to help students, parents, and educators with things like personalized guidance, communication, and data analysis.

Last year alone, PowerSchool says it blocked more than a billion web attacks in its work with K-12 districts. With the rapid development of artificial intelligence, technology leaders at the company knew they wanted to get ahead of anticipated challenges, take the initiative on strong security practices, and differentiate themselves from other education organizations that were also providing AI products.

"If you put something on the web, it's going to get attacked," said Mishka McCowan, vice president of cyberthreat management for the company. Twenty years ago, cyberattacks were relatively rare, but by a decade later they had become highly profitable for attackers, and now they've "blossomed into a multi-billion-dollar business," he said.

PowerSchool's first step in red teaming began with finding a company to do the work. There aren't many organizations with specialized expertise, so the company had to look for a security firm that was the right fit.

Among the questions they asked in screening vendors: What methodology do they use to test systems? What kind of professional background do the testers come from – if they are former web developers, PowerSchool wanted to know that they were capable of thinking with a cyberattacker's offensive mindset, rather than a protective, defensive one.

And were the red team companies subject matter experts on the products in question – in this case, PowerSchool's large language models?

The company PowerSchool eventually chose to perform the work was Cobalt Offensive Security Services, which has delivered about 15,000 manual security penetration tests to date. Its staff consists of members who wrote a commonly used standard for protecting large language models: the OWASP Top 10 for Large Language Model Applications.

The process for Cobalt Offensive Security Services began with a pre-test period, in which three testers were brought in, given login credentials, and briefed on the architecture of the system.

The clearer the security testers are on "how things work, the better results they can get without having to spend time on discovery," McCowan said. The goal was to be "collaborative" so that PowerSchool was giving the red team "information because we don't want them to waste time trying to figure it out," he added.

Then the testing period began. Over two weeks, the red teamers worked to find holes in the system.

"Nothing's off limits, they can do whatever they want to it," McCowan said. The goal in testing the defenses, he said, was clear: "They need to break it."

At the end of the process, red teamers came back and sat down with the company to go over the final report. During this time, developers had the opportunity to ask questions about what was exploited and how they did it.

"We work closely with our customers to support them through the remediation process, whether they need to update software or adjust some access controls," said Wong, Cobalt's chief strategy officer. "[We tell them,] 'Here's what we found that a bad person could do, and here's our recommendation on how to fix those things.'"

Few Standards, Low Expectations

The responsibility for data security falls on technology vendors, said Doug Levin, co-founder and national director of the cybersecurity nonprofit, K12 Security Information eXchange.

Most school systems do not assess the cybersecurity of companies seeking to work with them when they're considering products, he said.

That's partly because districts, with limited funds and resources, don't always have in-house expertise on cyberthreats, making it difficult for them to know what to ask for.

There are also few widely accepted indicators of trust in the K-12 sector when it comes to cybersecurity, Levin said, including any sort of "good housekeeping seal of approval."

"School systems are not routinely being held to a cybersecurity standard of practice, so it's not on their radar, and they haven't been asking about it during procurement," he said. "And because they haven't been asking about it during procurement, many companies have not felt like there's an incentive to invest in it."

Those weaknesses across the education sector create an opportunity for ed-tech companies that demonstrate initiative and transparency and take creative steps to protect their customers.

"Certainly, the notion that a company was regularly being tested and was willing to share its findings with their customers would make me more positively inclined toward them," Levin said.

Don Ringelestein shares that sentiment in his role as executive director of technology for Yorkville Community Unit School District 115, a district with 7,200 students in the suburbs of Chicago.

Cybersecurity is just not something that's usually top of mind for districts, he said. Although there are a handful of technology leaders who may come to the table knowing what questions to ask, most districts in Illinois don't have a chief information security officer, he added.

"People in my shoes would be a lot more confident if companies [went through red teaming]," he said. "We're sitting at the decision-making table. A red teaming exercise would be very valuable…for the vendors to be prepared to answer questions and to make sure problems are addressed prior to the purchasing of a system."

Post-Test Results

PowerSchool came away with two notable findings, as listed in their public report. The testers were able to manipulate prompts so that the AI assistant would change the topic. Students could have used that vulnerability to venture into topics that would otherwise be off-limits.

The red team review also found that certain prompts produced results of information the system uses to create responses. Although this wasn't a direct vulnerability, it would have allowed an attacker to examine what goes on behind the scenes in the platform to find other vulnerabilities.

In the last phase of PowerSchool's red team exercise, the company's internal teams took the findings and fixed the weaknesses, before arranging a retest, so that Cobalt could ensure that all vulnerabilities found were indeed remediated. All issues were fixed before the newest products were released, and the results of the test were compiled into a report that customers can access upon request.

The entire process from start to finish took about seven weeks. Cybersecurity experts say the length of the testing period can vary greatly, depending on the vulnerabilities that the red team finds.

The process was "an opportunity for us to learn and get better and incorporate that into other projects," said Rich Gay, chief information security officer at PowerSchool. "And customers have recognized the value of what we're doing."

School districts get the assurance that "we're not just saying we're doing these things," Gay added. "We're actually showing them what we found and [giving them] the confirmation."