TechnologyAdvice

July 2025

INDUSTRY TREND REPORT

Exposed and Exploited:

The Growing Risk of AI-Accelerated Vulnerabilities

SPONSORED BY

ThreatDown™
Powered by Malwarebytes

TABLE OF CONTENTS

Exposed and Exploited

The Growing Risk of AI-Accelerated Vulnerabilities

Any company serious about cybersecurity knows that the future of cybercrime will hinge on who can find vulnerabilities the fastest. For years, vulnerabilities have been at the heart of a tense race between attackers and defenders. Vendors, researchers, and bug bounty hunters search for weaknesses so they can fix them to keep businesses safe. Cybercriminals look for those same flaws so they can exploit them.

Both sides invest heavily in this search. Yet even when defenders gain an advantage by discovering a new vulnerability and issuing a patch, criminals can reverse engineer the update to discover what has been fixed, and attempt to exploit it. They do this because they know that many organizations will not install those fixes right away.

This delay between a patch being made available and it being applied, which can stretch over months, is known as the patch gap. During this window of time, attackers can figure out what has been fixed, and scan the internet for systems that have not yet been updated. The patch gap remains one of the most significant and stubborn challenges in cybersecurity.

Many security experts believe that artificial intelligence will transform how vulnerabilities are discovered and raise the stakes in the race to patch systems even further. Researchers have already shown that AI tools can analyze source code, detect patterns linked to security flaws, and even uncover zero-day vulnerabilities that no one knew existed.

Projects like Google's Big Sleep, research from the University of Illinois into automated zero-day discovery, and tools such as XBOW demonstrate that AI can handle complex tasks that once required significant time and human expertise. These innovations are intended to help defenders, but there is growing concern that cybercriminals might use similar technology to find weaknesses more quickly and on a larger scale.

Given the success of these projects, there is every reason to assume that attackers will also attempt to utilize AI for vulnerability research. Malicious versions of AI tools like XBOW and Big Sleep could help criminals identify flaws faster, expand their reach, and exploit systems during the patch gap with greater efficiency.

This shift is unlikely to create entirely new types of cyberattacks, but it could allow familiar tactics to move faster and affect more targets. That would put even greater pressure on businesses to strengthen their defenses and keep up with a rapidly evolving threat landscape.

Mid-sized businesses face particular risks in this environment. Many lack dedicated security teams, sophisticated monitoring tools, or the resources to maintain round-the-clock vigilance. Without strong defenses, they could become easy targets in a future where AI speeds up both the discovery and exploitation of vulnerabilities.

ThreatDown sees this as a critical moment for preparation. The company believes businesses need to understand why vulnerabilities matter so much, recognize the risks created by the patch gap, and understand the seismic effect that AI is having on vulnerability research and discovery. This report explores what could happen if attackers adopt these same techniques and offers practical guidance on how organizations can protect themselves in a world where threats may start moving much faster.

Abstract

As artificial intelligence becomes more powerful and easier to use, it raises a big question: **What role will AI play in future cyberattacks?** AI is reshaping how security researchers do their work. It's giving them new ways to analyze systems and spot weaknesses much faster than they could before.

These advances are helping defenders strengthen cybersecurity, but there's growing concern that criminals could use the same technology to launch attacks more quickly and on a much larger scale. Right now, there isn't much public evidence that this is happening widely, but experts warn that the window for staying ahead may be closing.

This report looks at how these changes could impact businesses, especially those without big security teams or advanced tools. It draws on insights from cybersecurity expert Mark Stockley, who shares his views on how AI could shift the balance between attackers and defenders. It also explores how criminals might put these tools to use and offers practical steps organizations can take to stay protected as cyber threats start to evolve more rapidly.

> AI is changing the pace, not the playbook. That shift puts pressure on defenders to spot threats earlier and respond faster, especially as automation becomes more common on both sides of the fight. "We don't think AI will change the way that attacks happen. We think that the tactics that criminals use for the time being will be the same. What we think is they will get quicker."
>
> **Mark Stockley,**
> *ThreatDown*

A Shifting Threat Landscape: AI's Expanding Role in Cybercrime

Focus

Artificial intelligence is starting to change how cyberattacks happen, even if those changes aren't always visible at first glance. So far, the strongest evidence of AI being used in cybercrime has been in areas like phishing emails, voice cloning, and deepfake videos, where researchers and security teams can clearly see the technology in action.

Reports from AI companies like OpenAI and Anthropic have also shown that criminals are using tools like ChatGPT to research targets, write and debug malware, and even run entire social media influence campaigns.

This isn't happening everywhere yet—**but the ingredients are falling into place.**

Security researchers have also shown that AI agents have the potential to break into computer systems, simulate the behavior of criminal hackers, or even run entire ransomware attacks. These developments suggest that AI-driven cyberattacks are not just a distant idea, but something that businesses need to prepare for.

Mark Stockley of ThreatDown points out that malware created with AI can be hard to detect, but that doesn't mean it isn't already being used. "There's nothing that an AI-generated malware does that a human-generated malware doesn't do," he says. "The difference is scale and speed."

This is why defenders need to stay alert. Attackers might not have changed their tactics completely yet, but the time it takes to develop and launch new threats is shrinking.

Expert Insight

Mark Stockley notes that AI-generated malware is difficult to detect, but he believes criminals are likely using it already as a coding assistant. Reports from organizations like OpenAI and Anthropic show evidence that attackers are using AI to write malicious code, fix bugs, and work more efficiently.

Stockley emphasizes that while AI hasn't created entirely new types of attacks, it's making existing tactics faster and easier to carry out. He adds that AI is starting to raise important questions about how defenders will need to adapt as these tools continue to advance.

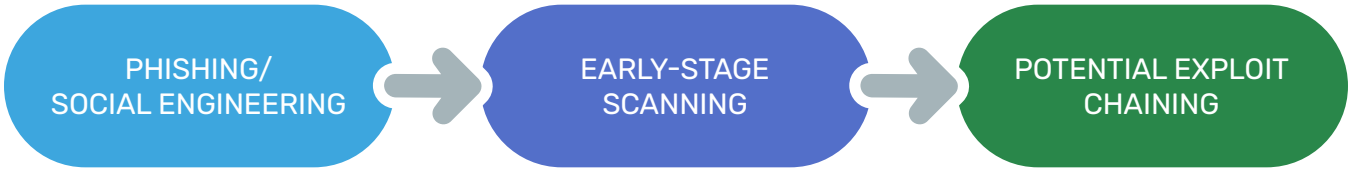

AI and the Future of Vulnerability Discovery — Lessons from XBOW and Big Sleep

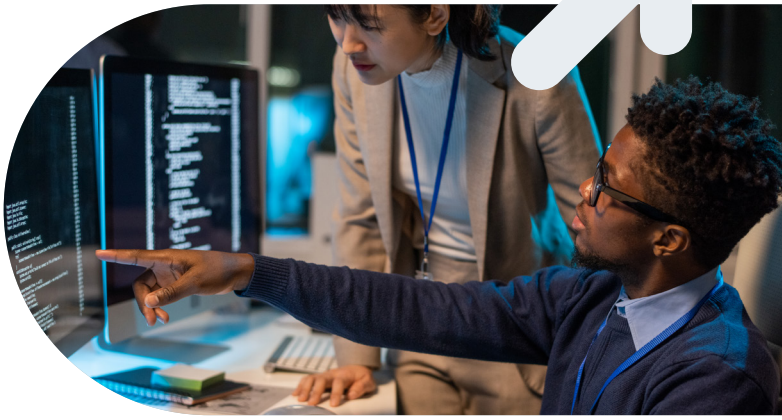

Focus

Much of the discussion around AI in cybercrime has focused on social engineering and phishing. However, some of the most significant breakthroughs in cybersecurity research involve using AI to discover vulnerabilities more quickly and at a larger scale.

Two projects blazing a trail in vulnerability discovery: Google's Big Sleep and a research tool called XBOW.

- **Google's Big Sleep**

  In 2024, Google Project Zero announced Big Sleep, an AI system designed to help identify vulnerabilities across large codebases. The system used deep learning and pattern recognition to spot subtle flaws that human researchers might overlook. Google reported that Big Sleep reduced the time required to discover certain vulnerabilities from weeks to just hours in some cases. It also became the first AI to uncover a previously unknown zero-day in a widely used piece of software, SQLite. While this tool was created to strengthen cybersecurity, it also shows how AI can accelerate the discovery of flaws that had gone unnoticed.

- **XBOW (eXploit Builder Without Human)**

  XBOW is a commercial platform that has quickly gained recognition as one of the most effective AI-based vulnerability researchers in the field. In 2025, it was named the top bug bounty hunter in the United States, outperforming human researchers by finding zero-day vulnerabilities, generating exploit code, and submitting high-quality reports that earned real-world bounties.

  The system automates the entire vulnerability discovery process, from detection to proof-of-concept creation to disclosure. Built initially to enhance defensive research, XBOW now shows that AI can already rival and, in some cases, surpass human capabilities in offensive security work. Its success raises concerns that similar tools could eventually be used by attackers as well.

These tools were developed to benefit defenders. Their existence shows that AI can now take over significant portions of vulnerability research and exploit development that were once manual and time-consuming. Security experts warn that similar technology could one day be used by cybercriminals to:

- Find far more vulnerabilities than before
- Reduce the time needed to develop exploits
- Target systems during the patch gap with greater efficiency

For now, the risk is not that AI will invent entirely new kinds of attacks but that it will dramatically increase the speed and scale of existing ones. Tools like Big Sleep and XBOW demonstrate the potential for AI to quickly identify weaknesses, which could enable attackers to outpace defenders who are still relying on traditional methods.

For businesses, this raises serious challenges:

- Security teams must anticipate that attackers could discover vulnerabilities faster than patches are currently being applied.
- Companies should work to reduce the patch gap through automation and rapid patch testing.
- Maintaining visibility into all assets is critical, because, as Mark Stockley has said, "You can't patch something if you don't see it."
- Organizations should consider adopting vulnerability scanning and management tools that incorporate AI capabilities to keep pace.

While Big Sleep and XBOW were created to protect systems, they also serve as proof that the tools helping defenders today could become blueprints for attackers tomorrow. Businesses must be ready for a future where cybercriminals might deploy AI-driven vulnerability discovery at speed and scale.

Why Mid-Sized Businesses Could Be Caught Off Guard

Focus

Mid-sized businesses often exist in a difficult middle ground. They're large enough to be attractive targets for attackers, but not large enough to maintain fully staffed security operations centers. Many have no dedicated security personnel or lack the resources for continuous system monitoring. Frequently, the same individuals who handle daily IT operations are also responsible for managing cybersecurity.

SMBs are being targeted nearly **four times more** than large organizations.

Source: Verizon

Mark Stockley explains that these organizations may have dozens of digital entry points, yet have just one or two people monitoring them, and often only during business hours. He notes that large enterprises typically have security teams working around the clock, while most mid-sized businesses do not. As a result, many IT teams leave systems unmonitored during the hours when cybercriminals are often most active.

Ransomware groups, for instance, commonly launch attacks in the early morning hours when IT staff are less likely to detect malicious activity. The use of AI could make such attacks far more common, increasing the strain on businesses already operating with limited resources.

Compounding the issue, many mid-sized organizations lack advanced tools for real-time monitoring and detection. Without sufficient visibility, minor incidents can escalate into significant breaches before they are identified or contained. For this reason, experts emphasize that mid-sized businesses require security solutions that integrate easily with existing systems and reduce operational complexity, helping lean teams maintain focus without becoming overwhelmed.

Expert Insight

Mark Stockley notes that AI-generated malware is difficult to detect, but he believes criminals are likely using it already as a coding assistant. Reports from organizations like OpenAI and Anthropic show evidence that attackers are using AI to write malicious code, fix bugs, and work more efficiently.

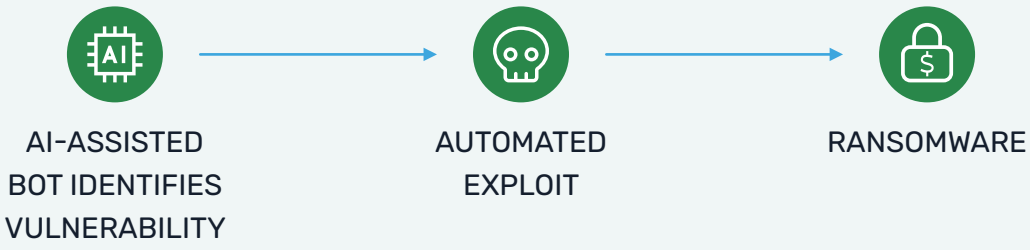

Beyond Patching: Why Speed and Layered Defense Matter

Focus

In cybersecurity, the advice to “patch early, patch often” remains sound guidance. However, Mark Stockley explains that while many organizations understand how important the guidance is, many also find it difficult to follow, due to operational demands, compatibility concerns, or limited resources.

Attackers exploit this delay. As soon as a patch is released, they analyze it to learn what vulnerability it fixed, create software to exploit it, and then search for systems that remain unpatched. This process can be extremely quick, and in rare cases, criminals have been able to begin exploiting vulnerable companies within hours of a patch being published.

The period between a patch being made available by vendors and actually being applied by businesses is known as the patch gap, and it can last for months, giving cybercriminals ample time to scan for a large number of exposed systems and launch attacks.

This is why automation and layered defenses are essential. Automating vulnerability assessment and patch management helps reduce the patch gap. Yet defenders also need to plan for cases where patching alone cannot prevent an attack, especially when facing zero-day vulnerabilities with no fix available.

Stockley compares attackers to burglars who move quietly through a house. The more obstacles they encounter, the more noise they make, giving defenders a better chance to detect suspicious activity before serious harm occurs.

For mid-sized businesses, adopting this mindset is crucial. Prevention alone cannot guarantee protection. The ability to slow down attackers, detect unusual behavior, isolate threats, and recover quickly can determine whether an incident stays minor or becomes a major crisis.

> "The first step to effective patching is making an inventory of your devices. **You can't patch a computer if you don't know it exists.**"
>
> **Mark Stockley**
> *ThreatDown*

Real-World Patch-Gap Exploits

When organizations delay applying fixes, attackers turn that window into a roadmap for invasion. Below are data points that tie breaches and exploit campaigns directly to unpatched vulnerabilities:

- **Log4Shell Vulnerability**

  The Log4Shell vulnerability was a major security flaw in Log4j, a tool used by countless applications for logging information. Although the bug existed quietly for years, it was discovered in 2021 and turned out to be incredibly dangerous. Hackers found they could send sneaky messages containing special code into apps that used Log4j, tricking those apps into reaching out to servers they controlled and downloading malicious software. This gave attackers a way to run commands on the victim's systems, sometimes taking full control. The flaw led to a scramble across the tech world, as companies rushed to patch millions of systems to stop hackers from exploiting it for data theft, ransomware, or other attacks.

- **Equifax Data Breach**

  In 2017, Equifax suffered one of history's biggest data breaches when hackers exploited an unpatched vulnerability in the company's Apache Struts software, ultimately stealing sensitive data from over 147 million people, including Social Security numbers and credit details. The breach remained undetected for months, leading to public outrage, multiple investigations, and a settlement of up to $700 million. It became a major wake-up call about the risks of delayed patching, poor security practices, and the need for stronger data protection. Today, the Equifax breach still shapes conversations about cybersecurity and how companies safeguard personal information.

Together, these findings confirm that patch gaps are not only common but also directly contribute to successful cyberattacks. As Mark Stockley notes in this report, attackers actively exploit the time between a patch's release and its deployment in real-world systems, creating a dangerous window of vulnerability.

How Fast Attackers Strike vs. Slow Patch Cycles

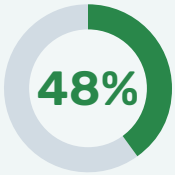

Proportion of new vulnerabilities weaponized in under 19 days

Source: Edgescan 2024 Vulnerability Statistics Report

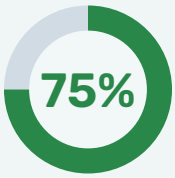

Organizations taking up to three weeks to apply critical patches

Source: Help Net Security – NVD Vulnerabilities

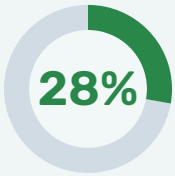

Percentage of vulnerabilities still unpatched after 12 months

Source: Axios (Synopsys survey)

30 days
(down from 60 days in 2022)

Average remediation time for critical flaws (Feb–Aug 2024)

Source: CISA Cybersecurity Performance Goals Adoption Report

5 days average

Time-to-exploit (TTE) for newly disclosed high-severity vulnerabilities

Source: Google Cloud Threat Intelligence: Time-to-Exploit Trends 2023

Expert Insight

Stockley emphasizes that attackers can reverse-engineer newly released patches within hours, creating a dangerous window of opportunity before organizations have time to respond. He advises that layered security is essential, explaining, "Patch management is about keeping attackers out. EDR is about catching them when they're already inside."

Getting Ahead of the Curve: How to Prepare Now

Focus

Cybercriminals have not completely changed how they operate overnight, but that could start to shift as agentic AI becomes more advanced. Unlike generative AI, which simply helps people write or code, agentic AI can carry out tasks on its own without constant human input.

Mark Stockley describes this shift as a major turning point that is already starting to unfold. “Agentic AI arrived in 2025. It’s still early days. But the entire AI industry is basically pivoting to agentic AI. That is the next big thing,” he says. He points out that until now, the scale of ransomware attacks has been limited by how many people are willing or able to participate. With agentic AI, cybercriminals could potentially automate entire operations, launching multiple attacks at once without needing large teams. For defenders, this means the pace and scale of threats could increase, making preparation even more important.

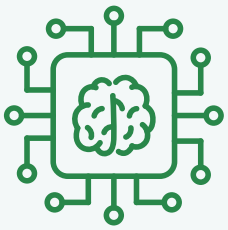

Agentic AI is the next big thing.

Businesses looking to stay ahead of these changes should focus on a few key strategies:

- Automating patching and vulnerability scans to reduce the openings attackers can exploit
- Using EDR or MDR tools to catch threats that slip past initial defenses
- Working under the assumption that someone might already have access to the network
- Choosing security platforms that integrate easily and keep noise and false alerts to a minimum

ThreatDown combines these capabilities into a single solution, tailored for mid-sized IT teams. Its platform handles automation, fast recovery, and constant monitoring so security teams can act quickly, even during off-hours. This means businesses do not have to hire extra staff or build complicated operations from scratch. Instead, they can rely on ThreatDown to manage the heavy lifting, making security more manageable and less overwhelming.

The biggest concern isn’t that cybercriminals will invent entirely new attacks. The real threat is that familiar attacks could happen much faster and on a much larger scale. In this environment, being prepared is not optional. It’s essential.

Expert Insight

Mark Stockley emphasizes that if AI makes attacks faster, defenders have to become faster too. He explains that security teams will need ways to slow attackers down and improve how quickly they can see and respond to threats. Automation, visibility into what's happening, and tools that can quickly reverse damage are all crucial.

Stockley also warns that agentic AI, which can act without human direction, might one day let attackers run entire ransomware campaigns with very little human involvement. In such a world, simple and effective tools that cut through noise and help teams respond quickly will be critical for staying secure.

Organizations interested in staying ahead of these new challenges can explore how **ThreatDown** offers solutions designed to help mid-sized businesses keep pace with the evolving threat landscape.

ThreatDown™
Powered by Malwarebytes

Born from Malwarebytes. Built for Business.

Formerly known as Malwarebytes for Business, ThreatDown delivers cybersecurity for businesses.

Built on Malwarebytes' expertise in malware detection and remediation, ThreatDown's comprehensive suite of solutions is designed for IT-constrained teams that need advanced protection, without the complexity.

Improve your security posture with one easy-to-manage platform that combines EDR, MDR, Patch Management, 7-day Ransomware Rollback, and AI-driven insights— deployed with a single, lightweight agent. No extra headcount or complex setup required.

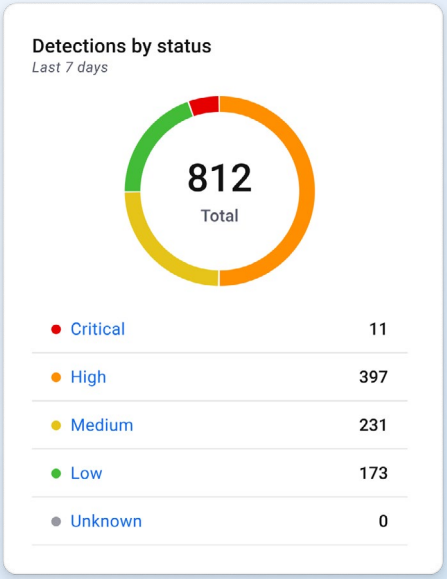

MRG Effitas CERTIFIED
ThreatDown Powered by Malwarebytes
PRODUCT OF THE YEAR 2025

SPRING 2025
Grid Leader

IT teams trust us to deliver enterprise-grade protection that's easy to deploy and manage. From award-winning tools to top ratings on Gartner Peer Insights™ and G2, security professionals choose ThreatDown for ease-of-use, fast deployment, and responsive support.

Spend less time managing threats and more time powering your business.

Visit threatdown.com/explore-portfolio

Studio A
by TechnologyAdvice

Custom Content Created by StudioA

Research-driven content that delivers actionable insights and empowers business decisions.

TechnologyAdvice

SPONSORED BY ThreatDown™ Powered by Malwarebytes