AI vs. Cybersecurity Stocks

What Anthropic's Claude Announcements Really Mean for NET, CRWD, RBRK, ZS, NTSK, S, and PANW

An investor's plain-English guide to seven cybersecurity companies, two AI announcements that spooked the market, and whether the fear is justified.

Introduction: Why Did Cybersecurity Stocks Drop?

In February 2026, Anthropic — the maker of an AI assistant called Claude — announced a tool called Claude Code Security. Within hours, shares of major cybersecurity companies fell sharply. Then, on April 7, 2026, Anthropic doubled down, announcing an even more powerful AI model called Claude Mythos Preview alongside a cybersecurity initiative called Project Glasswing. The April announcement triggered a second, larger selloff on April 10: CrowdStrike dropped around 8%, Cloudflare fell more than 13%, and Zscaler, Netskope, SentinelOne, and Palo Alto Networks all tumbled. Billions of dollars in market value were erased.

The fear was simple: if an AI can now find and fix security problems in software automatically, do we still need all these expensive cybersecurity companies?

This article is written for investors who want a clear answer to that question. We'll explain what each of these companies actually does, what these AI announcements actually are, and whether the fear is justified for each one. The answer varies significantly from company to company.

Part 1: A Plain-English Primer on Cybersecurity

What problem does cybersecurity solve?

Imagine a large company — a bank, a hospital, a retailer. That company has thousands of computers, servers, and software programs, all connected together and to the internet. It stores sensitive information: customer data, financial records, medical files, employee passwords.

Criminals, foreign governments, and hackers constantly try to break in. They might want to steal data, lock down systems and demand ransom (called ransomware), or simply cause disruption. The companies we'll discuss exist to stop those attacks.

Here is the critical nuance: cybersecurity is not one single thing. It is a collection of very different problems, each requiring its own solution. Think of it like protecting a house:

you need a front-door lock, window latches, an alarm system, a safe for valuables, and perhaps a security guard. Each does a different job. Replacing one does not replace the others.

> **Key Term: Vulnerability**
> A vulnerability is a weakness or flaw in software code that a hacker could exploit to break into a system. Think of it like a cracked window in a house — it is not broken yet, but it is an entry point waiting to be used.

The Three Layers of Cybersecurity

To understand what these companies do, it helps to think about security in three layers:

1. Prevention (before an attack): Finding weaknesses in your software and fixing them before hackers find them first. This is like inspecting your house for unlocked windows.

2. Detection and Response (during an attack): Monitoring your systems 24/7 to catch hackers in the act and stop them as fast as possible. This is your alarm system and security guard.

3. Recovery (after an attack): Making sure that even if hackers succeed, you can restore your data and get back to normal quickly. This is your safe and your insurance policy.

Different cybersecurity companies specialize in different layers. This is critical to understanding whether Anthropic's AI announcements threaten their business.

> **Key Term: Endpoint**
> An endpoint is any device connected to a network — a laptop, a smartphone, a server. 'Endpoint security' means protecting those individual devices from being hacked.

> **Key Term: Zero Trust**
> Traditional security assumed that anyone inside a company's network was trustworthy. Zero Trust is a newer philosophy that says: trust nobody by default, verify everyone every time — even employees. It is like requiring a badge scan to enter every room in a building, not just the front door.

> **Key Term: SASE (Secure Access Service Edge) — pronounced 'sassy'**
> A security model that combines network security (controlling who can connect to what) with cloud delivery (doing it all over the internet). It is designed for a world where employees work from anywhere.

> **Key Term: Ransomware**
> A type of cyberattack where criminals break into a company's systems and encrypt (lock) their data, making it unusable. They then demand a ransom payment to unlock it. Even large companies and hospitals have been crippled by ransomware attacks.

Part 2: What Each Company Does

CrowdStrike (CRWD) — The Security Guard

CrowdStrike is one of the most well-known names in cybersecurity. Its main product is called the Falcon platform, and its job is to sit on every device in a company — every laptop, server, and computer — and watch for signs of an attack happening in real time.

Think of CrowdStrike as a highly sophisticated security guard that never sleeps. It watches everything happening on your devices: which programs are running, what files are being opened, what connections are being made to the internet. When it sees something suspicious — a pattern that matches known hacker behavior — it raises an alarm or automatically blocks the threat.

The Falcon platform has expanded well beyond its original endpoint protection roots. It now covers endpoint security (protecting individual devices), identity protection (ensuring hackers cannot steal employee login credentials), cloud security (protecting a company's cloud resources), threat intelligence (studying how hackers operate), and a next-generation SIEM (a system that collects and analyzes security data from across an entire organization). CrowdStrike describes its mission simply: "We stop breaches."

CrowdStrike is fundamentally a runtime security company — meaning it protects you while your software is running in the real world, catching live attacks in progress. As of early 2026, CrowdStrike has been a recognized leader in the Gartner Magic Quadrant for Endpoint Protection Platforms for six consecutive years.

Zscaler (ZS) — The Traffic Inspector

Zscaler solves a problem that became critical as companies moved to remote work and cloud computing: how do you secure internet traffic when your employees are working from home, a hotel, or a coffee shop rather than a company office?

Traditional security worked by building a fortress around the company's physical office: a firewall at the front door that checked everything coming in and out. But when employees started working from everywhere and using cloud software, that fortress model stopped working. Traffic was no longer going through the front door.

Zscaler's Zero Trust Exchange is a cloud-based security checkpoint that all internet traffic passes through, no matter where the employee is. Every time someone in the

company tries to visit a website, access a cloud application, or send data, it goes through Zscaler's inspection first. Zscaler checks: Is this safe? Is this person allowed to do this? If yes, it passes through. If no, it is blocked. Zscaler operates this exchange across more than 150 data centers globally, making it the world's largest inline cloud security platform.

> **An analogy for Zscaler**
>
> Imagine every road in a city funneled through a single inspection station before cars could go anywhere. Zscaler is that inspection station for internet traffic. It does not matter if you are driving from home, a hotel, or an office — you still pass through the checkpoint.

Cloudflare (NET) — The Internet's Infrastructure Layer

Cloudflare is different from most cybersecurity companies in that it is, at its core, a network infrastructure company that also provides security services. It operates one of the largest networks in the world, handling traffic for a significant portion of the internet.

Cloudflare's primary jobs include DDoS protection (when hackers flood a website with millions of fake visitors to take it offline, Cloudflare absorbs and filters that traffic), a Web Application Firewall or WAF (a filter that inspects incoming web traffic and blocks known attack patterns), content delivery (storing copies of websites closer to users so they load faster), DNS services (the internet's phone book), and Zero Trust access control.

A crucial distinction: Cloudflare is protecting traffic flowing across the internet in real time. It is not analyzing your source code. Cloudflare has also been aggressively building AI infrastructure products — Workers AI, AI Gateway, and a Firewall for AI — positioning itself as essential infrastructure for the AI era.

Netskope (NTSK) — The Data Watchdog

Netskope is a specialized SASE company. Like Zscaler, it focuses on securing internet traffic and cloud application access. Its particular strength is data security — preventing sensitive information from leaking out of a company.

Here is the problem Netskope solves: employees use dozens of cloud applications — Google Drive, Slack, Salesforce, Dropbox, and more. What happens when an employee accidentally uploads a file containing customer credit card numbers to their personal Dropbox? Or copies sensitive client data into an unauthorized AI tool? Netskope watches all of that cloud activity and can detect and block sensitive data from going to the wrong places. This capability is called Data Loss Prevention, or DLP. Netskope is considered one of the top leaders in inline DLP, alongside Zscaler.

Netskope differentiates itself through its single-pass architecture, AI-driven inspection, and its NewEdge network. SASE is essentially the company's entire focus — which is both a competitive advantage and a concentration of risk.

> **Key Term: DLP (Data Loss Prevention)**
> Technology that monitors, detects, and blocks sensitive information from leaving a company's control without authorization. Think of it as a guard who checks every package leaving the building to make sure nobody is smuggling out confidential documents.

Palo Alto Networks (PANW) — The Security Platform Giant

Palo Alto Networks is one of the largest dedicated cybersecurity companies in the world. Unlike the more specialized companies above, Palo Alto has built a sprawling platform that tries to cover nearly every category of security in one integrated suite.

Its product lineup includes network security (next-generation firewalls), cloud security via its Prisma platform, and security operations via its Cortex platform. Palo Alto has been aggressive in acquiring smaller security companies and stitching them together. Its strategy is consolidation — convincing large enterprises to replace their patchwork of different security vendors with a single, unified platform. This makes their business stickier but also means they compete, at least partially, in nearly every security category.

Notably, Palo Alto CEO Nikesh Arora stated in an earnings call just days before the Claude announcements that AI will not replace cybersecurity products anytime soon. Palo Alto became a founding partner of Anthropic's Project Glasswing.

SentinelOne (S) — The AI-Native Defender

SentinelOne competes most directly with CrowdStrike. Its Singularity platform also sits on devices and monitors for attacks in real time. What makes SentinelOne distinctive is that it was built from the ground up with AI at its core, rather than adding AI features on top of older technology.

SentinelOne's approach is designed for autonomous response: the system can identify and contain a threat without waiting for a human analyst or a software update. Its Purple AI feature functions as an AI-powered security analyst embedded in the platform — it helps human security teams investigate incidents faster by automatically summarizing what happened, suggesting next steps, and executing pre-approved response workflows. Purple AI has been recognized as one of the fastest-growing products in the company's history. SentinelOne was named a Leader in the 2025 Gartner Magic Quadrant for Endpoint Protection Platforms for the fifth consecutive year.

Rubrik (RBRK) — The Data Recovery Specialist

Rubrik approaches cybersecurity from a completely different angle than the other companies on this list. While the others focus on preventing or detecting attacks, Rubrik focuses on the aftermath: what happens after an attack succeeds?

Rubrik's core product is data backup and recovery with a cybersecurity focus. It creates secure, immutable copies of all a company's data on a continuous basis. 'Immutable' means the backups cannot be modified or deleted — even by hackers who have broken into the system, or by a rogue employee. This is critical because many ransomware attacks specifically target backup systems to prevent recovery.

When a company gets hit by ransomware, Rubrik allows them to restore clean data from before the attack and get back to business without paying the ransom. Rubrik also helps companies understand what data was accessed or stolen during an attack. As of January 2026, Rubrik reported subscription ARR (annual recurring revenue) of $1.46 billion, growing 34% year over year — a business accelerating, not stalling.

Rubrik has also launched Rubrik Agent Cloud, which monitors and audits AI agent activity in enterprise environments to prevent AI-related data breaches. This positions them as infrastructure for safe AI deployment.

> **Key Term: Immutable Backup**
>
> A backup copy of data that cannot be altered or deleted by anyone once it is created — not by hackers, not by employees, not even by accident. It is the equivalent of a read-only copy stored in a vault. If ransomware encrypts all your live data, you restore from the immutable backup and carry on.

Part 3: What Did Anthropic Actually Announce?

A Quick Word About Anthropic and Claude

Anthropic is an AI safety company founded in 2021 by former members of OpenAI. It makes Claude — an AI assistant that competes with ChatGPT. Claude is a large language model (LLM), a type of AI trained on enormous amounts of text and code, giving it the ability to read, write, reason, and answer questions at a high level. As of early 2026, Anthropic reported a $30 billion annual revenue run rate, surpassing OpenAI on that metric and signaling its rapid rise as a top-tier AI company.

Announcement #1: Claude Code Security (February 20, 2026)

Claude Code is a specialized version of Claude designed for software developers. It can read entire codebases — meaning it can read and understand a company's computer code — and help developers write, fix, and improve software.

Claude Code Security is a feature added to Claude Code in February 2026. In plain English: you give it access to your software's source code (the actual written instructions that make your program work, before it is turned into a running application),

and it reads all of that code looking for security vulnerabilities: the cracks and weak points that hackers could exploit.

What makes it different from older scanning tools? Traditional code scanners look for known, predefined patterns — like a spell-checker that only knows a fixed list of misspellings. Claude Code Security reasons about code the way a human security expert would: it understands what the code is trying to do, how data flows through the program, and where something could go wrong in unexpected ways. According to Anthropic, Claude Opus 4.6 identified over 500 previously unknown high-severity vulnerabilities in production open-source codebases during internal testing.

> **Key Term: Source Code**
>
> The human-written instructions that make software work. Before you can run a program, source code must be compiled (translated) into a form the computer can execute. Claude Code Security analyzes the source code — before the program is ever run or deployed.

Announcement #2: Claude Mythos Preview and Project Glasswing (April 7, 2026)

This was the bigger, more alarming announcement. On April 7, 2026, Anthropic officially unveiled Claude Mythos Preview, described as its most powerful AI model ever built. Anthropic was blunt: Mythos had become so capable at finding and exploiting software vulnerabilities that the company decided it was too dangerous to release to the general public.

During internal testing, Claude Mythos Preview found vulnerabilities in every major operating system and every major web browser — including some flaws that had gone undetected for decades. In one case, it fully autonomously identified and exploited a 17-year-old remote code execution vulnerability in FreeBSD. It found a 27-year-old vulnerability in OpenBSD. These are the kinds of discoveries that would typically require a highly skilled human security researcher working for weeks.

Rather than release Mythos publicly, Anthropic launched Project Glasswing: a controlled initiative that gives Mythos to a select group of companies — including Amazon Web Services, Apple, Broadcom, Cisco, CrowdStrike, Google, JPMorgan Chase, the Linux Foundation, Microsoft, NVIDIA, and Palo Alto Networks — specifically to use it for defensive security work. Anthropic committed up to $100 million in usage credits for this effort, plus $4 million in donations to open-source security organizations.

The logic: get the world's critical software secured by defenders before similarly capable AI becomes available to attackers.

> **Key Term: Zero-Day Vulnerability**
>
> A security flaw that was previously unknown to the software's developers — meaning there is 'zero days' of warning before it could be exploited. Zero-days are particularly dangerous because no patch exists yet. Claude Mythos Preview found thousands of them.

What These Tools Do NOT Do

This is the most important section for investors, because the market appeared to panic without reading the fine print. Neither Claude Code Security nor Claude Mythos Preview (in its current, restricted form) does any of the following:

- Monitors live network traffic
- Protects devices from active attacks in real time
- Detects hackers already inside your network
- Blocks ransomware or responds to incidents as they unfold
- Backs up or recovers your data
- Controls who is allowed to access what systems
- Stops a DDoS attack

Claude Code Security is, in essence, a powerful pre-shipment code inspector. Claude Mythos Preview, in its current form, is being given only to a small group of trusted partners for defensive use. Neither is a deployed commercial product competing with the platforms these seven companies sell.

Bank of America analysts stated directly after the February announcement that the tool only poses a significant threat to pure-play code scanning companies — meaning niche companies whose entire business is scanning code, like JFrog (which fell 25% that day) and GitLab (which fell 8%). Those companies are not on this list.

Part 4: How Does It Affect Each Company?

Now we can give a clear-eyed assessment for each of the seven companies. The framework is simple: how much of the company's business actually overlaps with what Anthropic's announcements do?

CrowdStrike (CRWD)

Threat Level: Low

CrowdStrike's Falcon platform operates entirely in the runtime world — it watches live, running systems. Claude Code Security operates entirely in the pre-deployment world — it reads code before it is ever turned on. These two products do not compete.

CrowdStrike CEO George Kurtz made this case publicly after the February announcement, and even prompted Anthropic's own Claude to confirm it: Claude Code Security is not designed to replace CrowdStrike's Falcon platform. The two tools address completely different problems at completely different points in time. His

argument: “AI doesn’t eliminate the need for security. It increases it. If you want to build AI, you need GPUs. If you want to deploy AI, you need security.”

After the April 7 Project Glasswing announcement, CrowdStrike surged 6.2% — its best single-day gain in over six months. The market recognized that being a founding Glasswing partner reframes CrowdStrike as a collaborator in Anthropic’s AI security push, not a victim of it.

One legitimate longer-term question: as AI platforms expand their capabilities over time, could they eventually encroach on territory closer to CrowdStrike’s runtime detection business? It is worth monitoring. But that is a 2028 or 2029 question, not a 2026 one. CrowdStrike itself is actively building AI capabilities into Falcon, including new AI runtime protection, shadow AI discovery, and AI-powered data security launched at RSA 2026.

> **CRWD Verdict:** The February and April selloffs appear to be overreactions for CrowdStrike. The business operates in a fundamentally different security layer, and CrowdStrike is a Project Glasswing founding partner. AI growth increases the attack surface that Falcon protects — making this a tailwind, not a headwind.

Zscaler (ZS)

Threat Level: Low to Moderate

Like Cloudflare, Zscaler’s core business is inspecting live internet traffic in real time. The Zero Trust Exchange is a network security tool — it sits in the middle of internet connections and decides what gets through. Claude Code Security reads code files in a repository. These are completely different problems.

Where Zscaler has a slightly more legitimate medium-term concern is in its data security products, particularly its inline DLP capabilities. If AI reasoning eventually makes it easier for competitors or AI-native tools to offer DLP capabilities without the years of investment Zscaler has made, that could create pricing pressure over time. However, this is not an imminent threat from either the Claude Code Security or Mythos announcements specifically.

Zscaler has been actively integrating AI into its own platform. Wedbush analyst Dan Ives commented after the selloff that firms like Zscaler are well-positioned to integrate AI and expand their capabilities. The AI wave is not landing on Zscaler’s core business.

> **ZS Verdict:** Fear is largely misplaced for Zscaler’s core zero trust network business. A mild medium-term concern exists in DLP, but that is a gradual competitive dynamic, not a sudden displacement. The selloffs look like overreactions.

Netskope (NTSK)

Threat Level: Moderate

Netskope's situation is more nuanced. SASE is essentially the only thing they do. That focused model is their competitive advantage, but it also means they have fewer product lines to fall back on if one area faces pressure.

Netskope's core value proposition — securing cloud app traffic and preventing data loss — is fundamentally network-layer, real-time work. Far from what Claude Code Security does. The immediate threat from either announcement is low.

However, Netskope's DLP and DSPM (Data Security Posture Management) products are areas where AI tools could gradually reduce the pricing power of Netskope's offering over the next few years. If AI reasoning becomes good enough to classify sensitive data in real time at low cost, that could commoditize part of what makes Netskope special. This is a multi-year question, not a today question.

Netskope is also a relatively recent public company, which means it has less scale and fewer financial reserves to weather prolonged competitive pressure compared to CrowdStrike or Palo Alto Networks. This makes it somewhat more vulnerable to investor sentiment shifts, even on indirect threats.

> **NTSK Verdict:** The immediate threat from both announcements is low. However, Netskope has a legitimate medium-term sensitivity in its DLP and data security products. Investors should monitor how AI-native tools evolve in the cloud security space over the next two to three years.

Cloudflare (NET)

Threat Level: Very Low — Possibly a Beneficiary

Of all the companies on this list, Cloudflare's selloffs are arguably the hardest to justify from a business fundamentals perspective. Cloudflare's core product is network infrastructure — running traffic through its global network, filtering DDoS attacks, managing DNS, and delivering content. None of that touches source code scanning or the kind of vulnerability research that Anthropic is doing.

In fact, Cloudflare is one of the companies most likely to benefit from the broader AI wave. More AI means more internet traffic, more connected devices, and more AI agents making API calls — all of which need to flow through network infrastructure that needs to be fast and secure. Cloudflare sits right in the middle of all of that.

Cloudflare has been building its own AI infrastructure products aggressively. In Q1 2025, the company reported a 4,000% year-over-year increase in Workers AI inference requests and a 1,200% increase in AI Gateway requests. The AI boom is a tailwind for Cloudflare, not a headwind. Cloudflare's management guided full-year 2026 revenue of $2.79 to $2.80 billion, ahead of analyst consensus of $2.74 billion — strong numbers that do not reflect a business under pressure.

There is one long-term, indirect question: if AI tools help developers write more secure code before deployment, there will theoretically be fewer vulnerabilities for Cloudflare's WAF to catch at the network layer. But this is a very slow, indirect causal chain, and even perfectly written code still needs DDoS protection and network-layer security that has nothing to do with code vulnerabilities.

> **NET Verdict:** The selloffs look like pure sector contagion for Cloudflare. The company's business has minimal overlap with Anthropic's AI security announcements and significant AI tailwinds. The underlying business is growing strongly. This looks like one of the most misplaced reactions in the entire selloff.

Palo Alto Networks (PANW)

Threat Level: Low — Actively Partnering

Palo Alto Networks is the most diversified cybersecurity platform on this list, which makes it both less vulnerable to any specific disruption and also partially exposed to many different security segments. Its products span network firewalls, cloud security, endpoint protection, and security operations.

Neither Claude Code Security nor Claude Mythos Preview in its current, restricted form directly threatens any of Palo Alto's core products, which are focused on runtime security, network protection, and security operations — not pre-deployment code scanning.

More telling is what Palo Alto has done in response: it became a founding partner of Project Glasswing, and its CEO noted that "By prioritizing defensive access to these powerful capabilities, Anthropic is helping ensure that while intelligence is being weaponized, the defenders are the ones with the superior stack." JPMorgan named Palo Alto its top cybersecurity pick after the April 7 announcement, calling it a critical layer in the emerging AI security stack.

Palo Alto has also been integrating AI into its own Cortex platform for faster threat detection and automated investigation. This positions them as a beneficiary of AI capabilities rather than a casualty.

> **PANW Verdict:** Palo Alto's selloff looks like sector contagion. As a Glasswing founding partner, Palo Alto is actively using Claude Mythos — the very model that spooked markets — to strengthen its own platform. JPMorgan named it their top cybersecurity pick after the announcement.

SentinelOne (S)

Threat Level: Low

SentinelOne is CrowdStrike's closest competitor — both are endpoint security and threat detection companies that watch live systems for attacks. Like CrowdStrike, SentinelOne's Singularity platform operates entirely in the runtime world. Claude Code Security works in the pre-deployment, source-code world.

SentinelOne's AI-native foundation is arguably an advantage here. The company was built around machine learning from the start. Its Purple AI platform — which uses generative AI to help security analysts investigate threats faster — is already doing what many companies are scrambling to add. SentinelOne was named a Leader in the 2025 Gartner Magic Quadrant for Endpoint Protection Platforms for the fifth consecutive year.

If anything, more capable AI tools raise the bar on attacker sophistication, which increases demand for platforms like Singularity that can respond at machine speed. The AI wave is a tailwind for a company built natively on AI.

> **S Verdict:** Low threat from both Anthropic announcements. SentinelOne operates in a different security layer (runtime endpoint protection) and its AI-native architecture positions it well for the broader AI transformation of security.

Rubrik (RBRK)

Threat Level: Very Low — Possibly a Beneficiary

Rubrik is arguably the least threatened company on this entire list. Its business — immutable data backup and ransomware recovery — addresses what happens after an attack succeeds. Claude Code Security addresses preventing certain types of vulnerabilities before software is deployed. These exist in entirely different universes.

You cannot AI your way out of needing a recovery plan. Even if every company in the world used AI tools to scan their code perfectly, hackers would still find ways in through other methods: phishing emails, stolen credentials, supply chain attacks, zero-day exploits. When those attacks succeed — and some always will — companies need Rubrik's recovery capabilities.

In fact, Anthropic's own warnings about Claude Mythos make the case for Rubrik stronger. If AI tools are going to make attacks more sophisticated and harder to defend against, the consequences of a successful breach become more severe — and the value of bulletproof recovery capabilities becomes greater. Rubrik's most recent financial results underscore this: subscription ARR reached $1.46 billion as of January 2026, up 34% year over year, with revenue growing 46% year over year in Q4.

Rubrik is also building AI-native capabilities of its own — including Rubrik Agent Cloud, which monitors and audits AI agent actions in enterprise environments. This positions them as infrastructure for safe AI deployment, not a casualty of it.

> **RBRK Verdict:** The most misplaced fear of all seven companies. Rubrik's recovery-focused business is entirely unaffected by a code scanner and may be strengthened by a more dangerous threat landscape. The selloff was pure sector contagion.

Part 5: The Big Picture for Investors

Who the Announcements Actually Threaten

To be direct: the companies with the most legitimate reason to worry about Claude Code Security are not the ones on this list. The real impact falls on pure-play code-scanning companies:

- JFrog: Fell 25% on the February announcement. Its core product is managing software builds and scanning code for security issues — directly overlapping with Claude Code Security.
- GitLab: Fell 8%. GitLab has a security scanning product embedded in its developer platform that competes more directly.
- Tenable: A vulnerability scanning and management company. More exposed than the network security names.

The seven companies on this list — NET, CRWD, RBRK, ZS, NTSK, S, PANW — are predominantly network security, endpoint security, identity, and data recovery companies. They do not primarily scan source code. The market painted them all with the same brush.

Why the Market Overreacted

Market selloffs driven by AI announcements tend to follow a pattern: investors hear 'AI does security thing' and sell everything with 'security' in the description, regardless of whether the specific AI capability competes with that company's actual products.

Raymond James analyst Mark Cash called the reaction 'excessive' and said investors were 'extrapolating well beyond the current functionality' of the tools. Robert W. Baird's Shrenik Kothari called it a 'panic-driven, narrative-led selloff.' BTIG noted the fears were 'based on an incorrect reading' of the announcements' impact. Bank of America stated the tool only poses a significant threat to code scanning platforms.

The Counterintuitive Reality: AI Is a Tailwind for Most of These Companies

Here is the argument that most market participants seem to be missing: more AI means more security spending, not less.

- More AI agents running on corporate networks means more endpoints to protect (good for CRWD, S).
- More AI-driven internet traffic means more network security demand (good for NET, ZS, NTSK, PANW).
- More sophisticated AI-powered cyberattacks mean more urgent need for recovery capabilities (good for RBRK).
- AI making attackers more capable means companies spend more on defense, not less.

Anthropic acknowledged this dual reality in its Claude Mythos announcement: the same capabilities that help defenders find vulnerabilities will also be used by attackers to find new ones. The cybersecurity industry is not threatened by AI — it is being given a faster, smarter opponent that requires a faster, smarter defense. That is a revenue opportunity for these companies, not a revenue threat.

One Legitimate Long-Term Question

None of this means these companies have zero exposure to AI disruption over a longer time horizon. As AI platforms like Claude expand their capabilities, they may gradually move into adjacent security functions. An AI that scans code today might monitor running systems or respond to incidents in two or three years. If Anthropic or other AI companies bundle increasingly powerful security capabilities into their platforms, that could eventually create budget competition with dedicated security vendors.

This is worth watching. But it is a multi-year story — and the companies on this list are not standing still. They are all actively integrating AI, and several (CrowdStrike, Palo Alto Networks) are already Anthropic partners.

Conclusion: Quick Reference Summary

Here is a simple summary of the assessment for each company:

Ticker	What They Do	Threat Level	Our Take
NET	Network infrastructure, DDoS, WAF, CDN, Zero Trust	Very Low — Beneficiary	Sector contagion. AI is a tailwind. Core business has zero overlap with code scanning. Strong revenue guidance for 2026.
CRWD	Endpoint & runtime security, identity, cloud, SIEM	Low	Overreaction. Different security layer entirely. Glasswing founding partner. Surged 6.2% on April 7.

Ticker	What They Do	Threat Level	Our Take
RBRK	Data backup, immutable recovery, ransomware defense	Very Low — Beneficiary	Most misplaced fear. Recovery-layer business unaffected. ARR $1.46B, up 34% YoY. AI threat wave may strengthen demand.
ZS	Zero Trust network traffic inspection, DLP	Low to Moderate	Core network business safe. Medium-term DLP competition worth monitoring. Oversold on current announcements.
NTSK	SASE, cloud security, DLP, data protection	Moderate	Narrow focus creates more vulnerability to DLP disruption long-term. Watch this space over next 2–3 years.
S	AI-native endpoint security, autonomous response	Low	AI-native platform is an advantage, not a liability. Runtime security unaffected. Named Gartner Leader 5 years running.
PANW	Broad platform: network, cloud, endpoint, SecOps	Low — Partnering	Glasswing founding partner. Using Claude Mythos for its own platform. JPMorgan top cybersecurity pick post-Glasswing.

The cybersecurity sector selloffs triggered by Anthropic's February and April 2026 announcements were driven largely by fear and narrative rather than a careful analysis of competitive overlap. For most of these seven companies, the actual business impact is minimal. In some cases, the AI revolution is actively creating more demand for their services.

That said, the AI transformation of security is real and accelerating. Investors should monitor how AI capabilities expand over the next two to three years and whether any of these companies' specific product lines face genuine pricing pressure or substitution risk. For now, for most of these names, the evidence suggests the market reaction was a buying opportunity for long-term investors — not a fundamental change in business outlook.
