J.P.Morgan



EYE ON THE MARKET | JULY 2026

Patchmageddon

The race to patch software vulnerabilities before zero-day cyber-exploitations proliferate

By **Michael Cembalest** | Chairman of Market and Investment Strategy for J.P. Morgan Asset & Wealth Management

Patchmageddon

The race to patch software vulnerabilities before zero-day cyber-exploitations proliferate[1]

Executive Summary. Imagine a world without tornado warnings. That was the situation that prevailed in the US until the late 1940's when the US Weather Bureau began issuing public tornado warnings, after which tornado deaths per million declined by 90%. Even though the average tornado warning lead time is just 15 minutes, that's enough for people to mitigate substantial risks to their survival.

A tornado may be coming in the form of previously unknown vulnerabilities in software and operational hardware that new frontier models can now detect. State-grade cyber capabilities will become increasingly accessible to actors with malicious intent such as ransomware operators, terrorists or hacktivists with no goal but destruction. The parallel to tornado warnings will come in the form of disclosed vulnerabilities and in the form of patches required to fix them. Many patches will need to be implemented as soon as possible since the average time between the disclosure of a vulnerability and its first exploitation has fallen to a single day (a zero-day event), leaving companies little more time to react than residents of tornado alley.

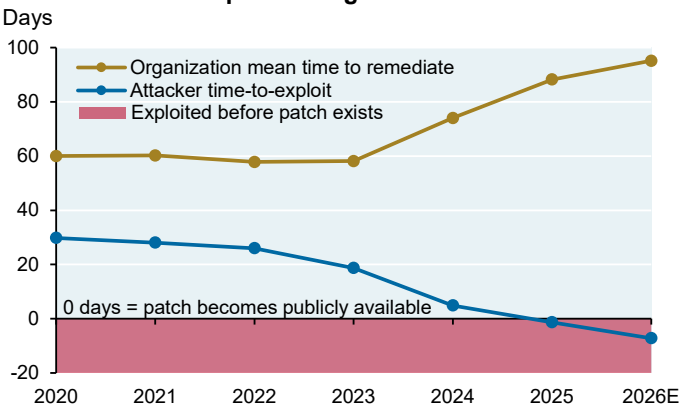
Unfortunately, it's taking longer for organizations to patch these vulnerabilities just as attacker time to exploit is falling. In addition, the rate of disclosed vulnerabilities was rising faster than the rate at which they're patched even before Mythos-grade models came along. Even when patches exist, companies often don't respond in time: in ~60% of breaches, a patch was already available at the time of compromise.

For software developers and maintainers, the most important epiphany is the following: the same tools that are used to detect and exploit vulnerabilities can also be used to propose code fixes and remediate them. Claude Security, GPT 5.5-Cyber and OpenAI's Codex Security Plugin are notable examples.

In this special issue Eye on the Market we review the world before Mythos and similar frontier AI models, how cyber risks have changed, the underappreciated breadth and risks from open source code, the proliferation of cheaper open weight models, the risks to physical infrastructure and some guidance on what business owners, software developers and the government should be doing to mitigate the potential consequences.

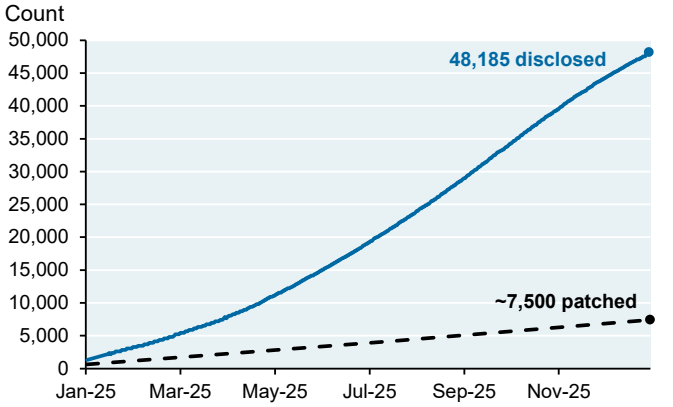
Michael Cembalest
JP Morgan Asset Management

Attacker time-to-exploit vs organization time to remediate



Source: Sidero Labs, May 4, 2026

Total common vulnerabilities disclosed vs patched, 2025



Source: Mitiga.io Cybersecurity, June 5, 2026

[1] I would like to acknowledge the substantial contributions to this piece from some of my JP Morgan colleagues: Pat Opet (Global Chief Information Security Officer), Brett Wallace (Global Head of Cyber Security Operations), Rusty Clark (Global Head of Cyber Security Intelligence), Robert Underwood (Chief Open Source Officer), Katheryn Rosen (Global Head of Regional Information Security), Terah Lyons (Global Head of AI and Data Policy), Adriana Villasenor (Head of Cybersecurity Communications) and Matt Robinson (Corporate Strategy)

Table of Contents

Glossary: the most important words and phrases to understand 3
The state of cyber risk before 2026 4
What has changed due to Mythos, GPT 5.5 and other models capable of enhanced detection capabilities 6
Chinese open weight models with advanced cyber capabilities and the jailbreak problem 10
Zero day risks to open source code and its army of volunteers 12
Zero day risks to physical infrastructure and the challenge of unpatchable legacy networks 18
For software developers and maintainers: tools that find vulnerabilities can also remediate them 20
The most important things for business owners to do 22
The most important things that the US government can be doing 23
Appendix A: on cyberattacks, the US and China 25
Appendix B: an example of open source dependencies: Express 4.17 26

Patchmageddon



Glossary: the most important words and phrases to understand

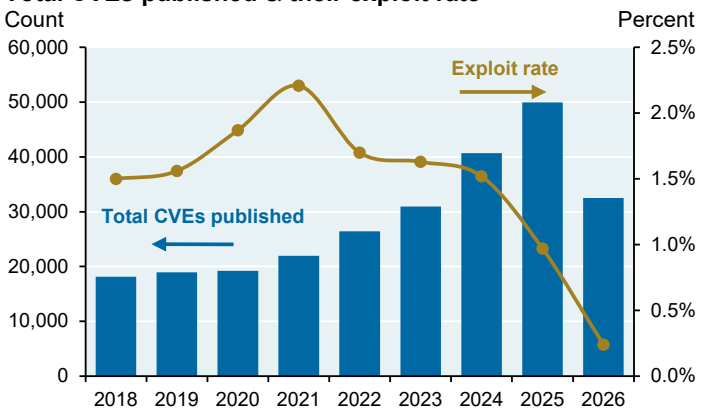
- **Abliteration:** a technique that removes built-in safety restrictions from models without retraining
- **Advisories:** formal notifications published by government agencies, security firms, or vendors detailing newly identified digital threats, vulnerabilities or malicious campaigns
- **Botnet:** a network of internet-connected devices like computers, smartphones and smart home gadgets that have been secretly infected with malware and are remotely controlled by a single cybercriminal
- **Common Vulnerabilities and Exposures (CVE):** standardized global database of publicly disclosed cybersecurity vulnerabilities and flaws in software and hardware
- **Cybersecurity and Infrastructure Security Agency (CISA):** part of DHS and lead federal agency for cyber defense and critical infrastructure protection
- **Exfiltration:** the unauthorized, intentional transfer or theft of data from a computer, device, or network to an external location controlled by an attacker
- **Fork:** a new, separate codebase that is created by duplicating an existing codebase, which allows users to iterate on ideas or changes to the code before proposing those changes to the original codebase
- **Github:** a developer platform that allows users to create, store, manage and share their code, most commonly used for open source software development projects
- **Jailbreak:** when hackers exploit vulnerabilities in AI systems to perform restricted actions by bypassing their built-in safety guardrails and ethical constraints
- **Maintainers:** the individuals who voluntarily manage and update open source code repositories by reviewing code, fixing bugs, releasing updates and patching security vulnerabilities
- **Open Weight vs Open Source**: Open weight means that the model and its parameters can be downloaded but its training code and dataset remain hidden. Open source means that a user can get full access to the model, the source code, its training data and architecture, allowing for independent recreation of the model itself
- **Patch:** a targeted fix for a specific issue or security vulnerability
- **Penetration testing:** an authorized attempt to break into a computer system, network or application to find security weaknesses before real attackers do
- **Programmable logic controllers (PLCs):** industrial computers used to control and monitor industrial equipment based on custom programming
- **Privilege escalation:** a cybersecurity technique where a threat actor exploits software vulnerabilities, system misconfigurations or design flaws to gain unauthorized, elevated access to resources
- **Red teaming:** method of testing cybersecurity effectiveness by simulating adversarial attacks against an organization's systems, strategies or defenses
- **Time-to-exploit (TTE):** measures the gap between CVE public disclosure and first confirmed in-the-wild exploitation
- **Vibe coding:** the practice of prompting AI tools to generate code rather than writing code manually
- **Zero-day exploitation:** a cyberattack that targets a previously unknown software or hardware vulnerability, giving the developer "zero days" to create and release a patch before the system is compromised

The state of cyber risk before 2026

The software industry's longstanding convention is to publicly disclose vulnerabilities 90 days after they're discovered (or if a patch is created before the 90 days is up, around 45 days after the patch becomes available). This delay is designed to allow time for users who have already been privately made aware of the problem to update or disable software before the vulnerability can be exploited by attackers.

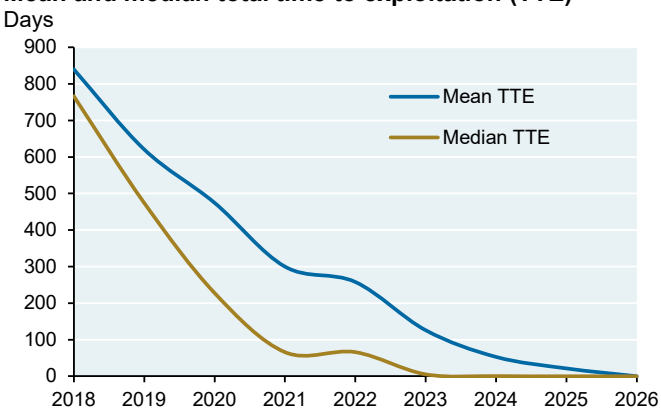
The challenge: even before Mythos and GPT 5.5 came along in 2026, companies and governments had to react with increasing speed once a vulnerability is disclosed. As shown in the first chart, the number of reported vulnerabilities has been rising, of which 1.5% -2.0% get exploited each year. And as shown in the second chart, the time between initial publication of a given vulnerability and its first confirmed exploitation (TTE) declined to less than 100 days on average in 2025, with median times to exploitation already falling to zero days. In fact, some attacks are occurring even before patches become available.

Total CVEs published & their exploit rate



Source: Zero Day Clock, July 2026

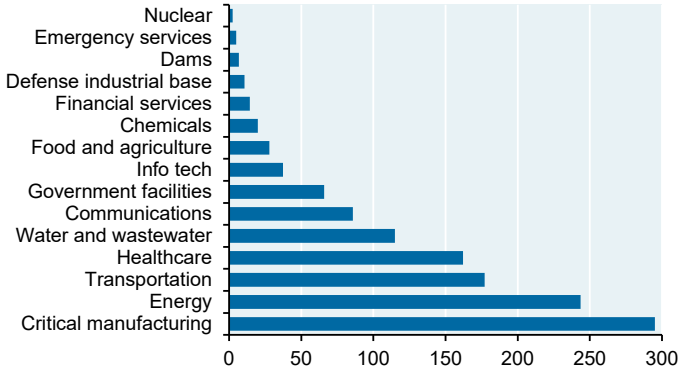
Mean and median total time to exploitation (TTE)



Source: Zero Day Clock, July 2026

Who gets targeted? Just about every sector although critical manufacturing[2], energy, transport and healthcare receive cyber advisories with a higher frequency according to advisories published by organizations like the Cybersecurity and Infrastructure Security Agency (CISA).

Cybersecurity vulnerability advisories by industry
Count of advisories, 2025



Source: Forescout Technologies, February 19, 2026

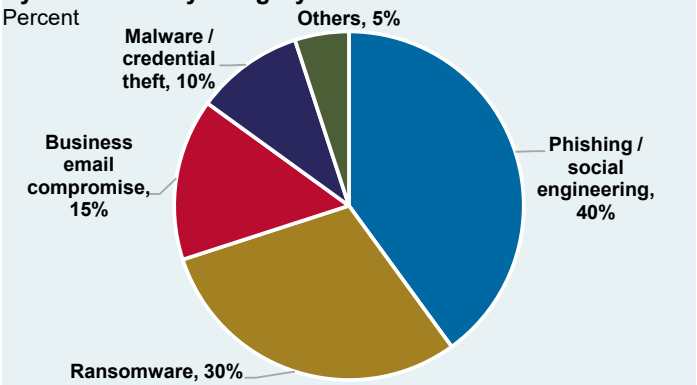
The **cyber advisories** in the chart are published by the Industrial Controls Systems (ICS) Advisory, a security bulletin published by CISA alerting organizations to vulnerabilities and threats affecting operational technology products. CISA has been the authoritative source on such vulnerabilities since they started the ICS Advisory program in 2010. Between March 2010 and January 2026, CISA published 3,637 ICS advisories on 12,174 vulnerabilities affecting 2,783 products from 689 vendors. Last year was also the first time that CISA published more than 500 ICS advisories in a single year

[2] Critical manufacturing is also the most frequent target of ransomware perpetrators. Improved cybersecurity practices, enhanced recovery capabilities and increased law enforcement actions contribute to declining profitability for ransomware operations. In 2025, ransomware payments amounted to ~$820 million, a 34% decrease from US$1.25 billion in ransom payments in 2023, although they have been increasing in healthcare. This decline has ultimately prompted attackers to shift toward data exfiltration instead of encryption. Source: *"Ransomware Threat Landscape: Q1 2026",* JP Morgan Chase Cybersecurity Operations, April 28, 2026

Some updates on the latest cyberattack trends:

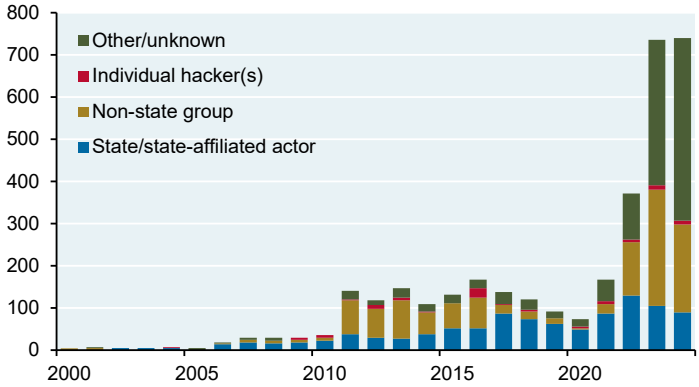
- Cyberattacks increased 18% globally in 2025 with 75,000 occurring every hour[3], and the average organization is attacked 2,000 times each week[4]
- Organizations with 1,000 – 10,000 employees were more than 2x as likely to report an incident than organizations with fewer than 100 employees
- Phishing remains a top attack vector with 33% of employees likely to be duped by phishing simulations[5]
- The global cybersecurity workforce is short approximately 4.8 million people[6]

Cyber attacks by category
Percent



Source: SQ Digital & Financial Magazine, April 2026

Cyberattacks over time by threat actor type
Number of attacks



Source: ECB Financial Stability Review, May 2025

A lot of cyberhacking activity is based in or financed by **China**. The global cyber threat landscape has evolved due to a strategic pivot by Chinese state-sponsored threat actors away from a primary focus on traditional espionage and intellectual property theft towards large-scale pre-positioning of malicious capabilities within the world's most critical infrastructure networks. China-nexus threat actors are increasingly using covert networks made up of compromised Small Office Home Office routers as well as Internet of Things and smart devices[7]. A Chinese state-sponsored cyber group known as Volt Typhoon compromised and maintained access within US energy, water, communications and transportation networks for as long as five years. And in September 2025, Chinese state-sponsored cyber actors targeted telecommunications, government, transportation, lodging and military infrastructure networks, stealing their data and tracking the movements of each network's users[8]. See Appendix A for research on Chinese cyberhacking of US and global entities.

Russia is also a major source of cybersecurity risk. Operation Endgame, the multinational law enforcement cyber operation targeting botnets and "cybercrime-as-a-service" infrastructure, announced in 2026 it had disrupted the Russia-domiciled SocGholish botnet. This network uses compromised WordPress sites to spread malware that purports to be software updates but instead provides criminals with access to victim computers and is often used for follow-on ransomware attacks[9].

[3] *"Cybersecurity attacks statistics 2026: rising threats now",* SQ Magazine, April 1, 2026

[4] "*Cyber security awareness month: ten things to know in 2025*", World Economic Forum, September 30, 2025

[5] *"2026 KnowBe4: Phishing by industry benchmarking report",* KnowBe4, July 7, 2026

[6] *"The State of the Cybersecurity Workforce",* ISC2, October 23, 2024

[7] "*Defending Against China-Nexus Covert Networks of Compromised Devices*", CISA, June 2026

[8] *"Countering Chinese State-Sponsored Actors Compromise of Networks Worldwide to Feed Global Espionage System",* CISA, September 3, 2025

[9] *"Open weight model advances make the Mythos debate moot",* Tom Uren, Lawfare, June 26, 2026

What has changed due to Mythos, GPT 5.5 and other models capable of enhanced detection capabilities

The cyber risk landscape has changed in a seismic way with the release of Mythos, GPT 5.5 and other models capable of detecting previously unknown vulnerabilities at large scale. We first wrote about this in April 2026 and highlighted some stark developments cited in Anthropic's Project Glasswing May 2026 update:

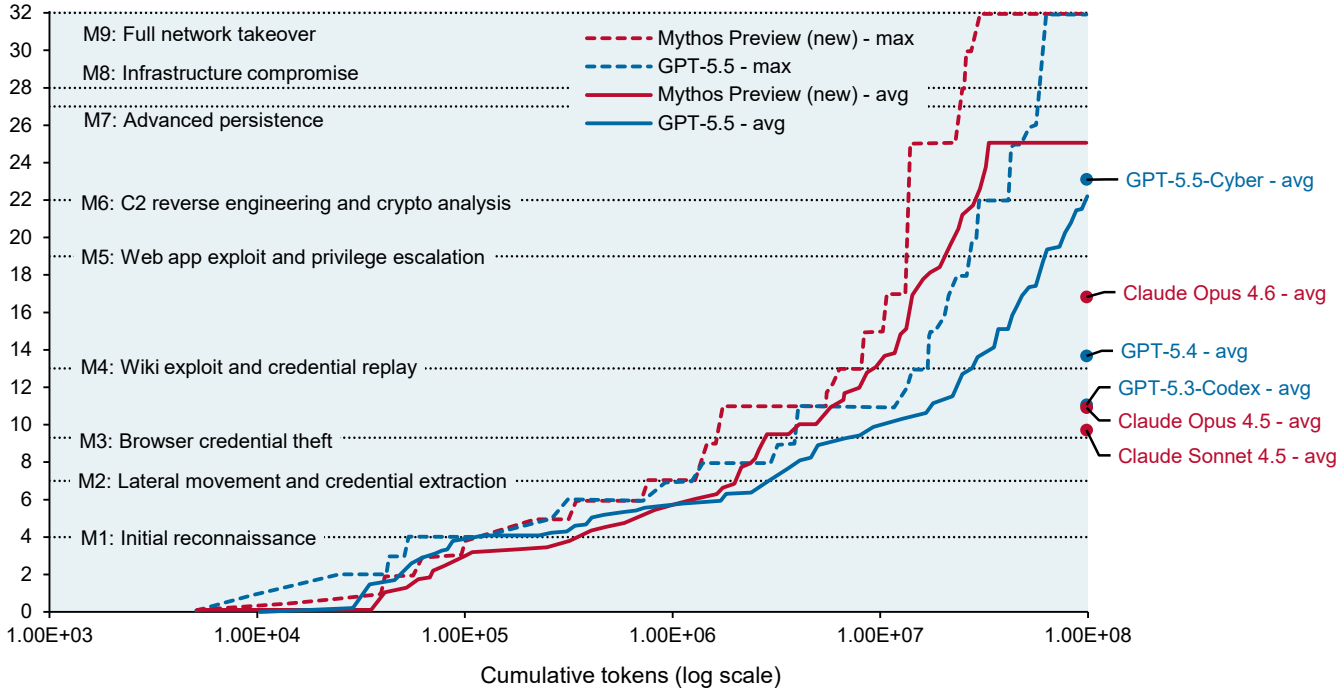
- Mozilla found and fixed 271 vulnerabilities in Firefox 150 while testing Mythos Preview, over 10x more than they found in Firefox 148 when using Claude Opus 4.6
- Anthropic and Project Glasswing participants found 10,000+ new high- and critical-severity zero-day vulnerabilities in its first month of use; Mythos also constructed a working certificate forgery exploit in a library used by billions of devices
- Cloudflare found 2,000 bugs, 400 of which were high or critical severity in the first month of use

The problem: 95% of the resulting Mythos vulnerability disclosures had no public advisory at the time of the research snapshot, meaning they were not visible through standard Common Vulnerabilities and Exposure reports, National Vulnerability Databases or GitHub advisory databases.

OpenAI's latest models have enhanced cyber detection capabilities as well. OpenAI launched the Daybreak Initiative to embed advanced AI models directly into software security and development workflows to help defenders detect, validate and patch vulnerabilities faster. OpenAI is working with a cybersecurity firm which cites a "firehose of security findings" and stated that its engineers worked with 19 open source projects using OpenAI's Codex and GPT-5.5-Cyber models to detect hundreds of legitimate bugs.

Another look at the heightened vulnerability identification capabilities of these new models comes from the UK-based AI Security Institute which evaluated Mythos and GPT 5.5 on cybersecurity tests. **Both Mythos and GPT 5.5. performed exceptionally well on a 32-step corporate network attack simulation spanning from initial reconnaissance through to full network takeover.** Before the latest models were released, frontier models were only able to get about halfway to total system takeover.

Completed steps per spent tokens on a 32-step corporate network attack simulation
Number of steps completed



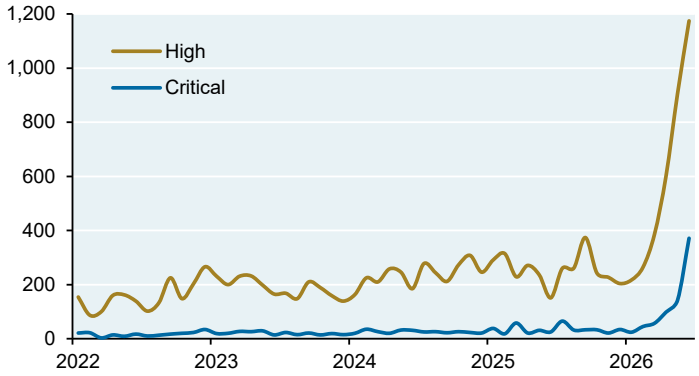
Source: AI Security Institute, May 13, 2026

More signs of the Pandora's Box that has been opened:

- Over the next 12-24 months, capabilities that are currently confined to a small number of nation-state cyber units are likely to be reproducible in freely available open weight models
- Models now have both the speed and ability to connect low- and medium-severity vulnerabilities into critical exploit paths including approaches no one has seen before. Writing an exploit from a known vulnerability and its corresponding patch is a largely mechanical process. Anthropic's own red team blog demonstrates this explicitly: Mythos Preview turned known vulnerabilities into working privilege escalation exploits in under a day at a cost of under $2,000 with zero human intervention
- The length and complexity of tasks an AI model can carry out independently is now doubling every 3- to 4-months (accelerating from every seven months before 2024)[10]

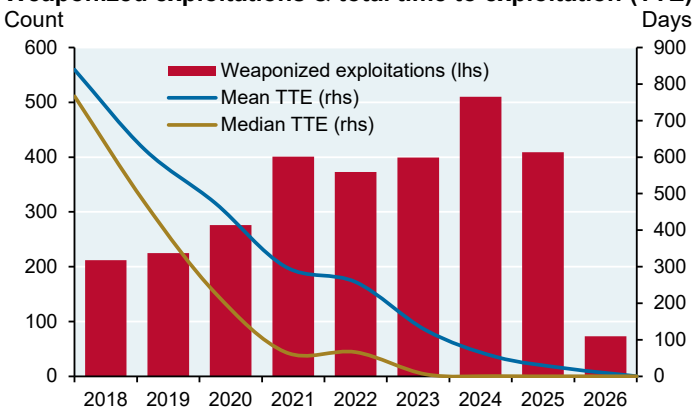
As shown on the left, the number of critical and high severity vulnerabilities reported by 21 tech companies has been surging this year[11]. While the number that have been exploited looks low so far according to the Zero Day Clock[12], there are already examples of novel exploitations taking place[13]. Note as well that mean and median time to exploitation has fallen to zero days.

Critical & high severity cybersecurity vulnerabilities (CVEs) from 21 notable organizations, Number of CVEs



Source: Epoch AI, July 2, 2026

Weaponized exploitations & total time to exploitation (TTE)



Source: Zero Day Clock, July 2026

[10] *"Measuring AI Ability to Complete Long Tasks"*, METR, March 19, 2025; *"Task-Completion Time Horizons of Frontier AI Models"*, METR, May 8, 2026

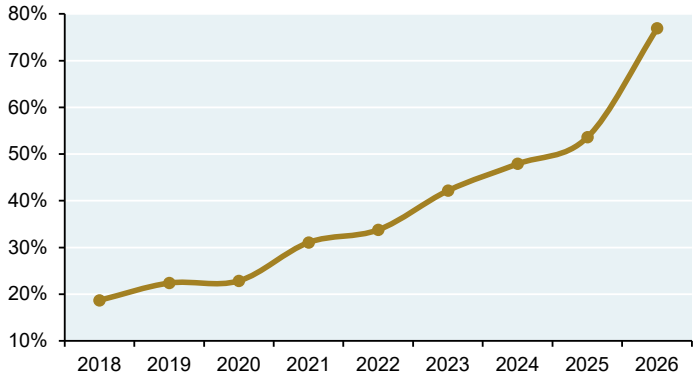
[11] The companies include: AWS, Apache, Apple, Cisco, Google, Linux, Microsoft, Mozilla, NVIDIA, Oracle, Red Hat, Adobe, IBM, Intel, AMD, Qualcomm, Samsung, SAP, VMware, GitHub and OpenSSL

[12] The **Zero Day Clock** is a live dashboard maintained by the Chief Information Security Officer of Sysdig, a cloud security company that assesses live threats and vulnerabilities across cloud environments

[13] Two examples from Mitiga.io: (a) the TanStack / GitHub Actions in May 2026 chained three flaws in a novel pattern: 84 malicious packages published in six minutes, an 18-minute window that quickly moved to GitHub, OpenAI, Mistral, and Grafana. Patching wasn't even an option; the gap was structural. (b) in May 2026, Google's threat intelligence group attributed the first publicly confirmed in-the-wild zero-day exploit developed with AI assistance, a 2FA bypass against a popular open source admin tool

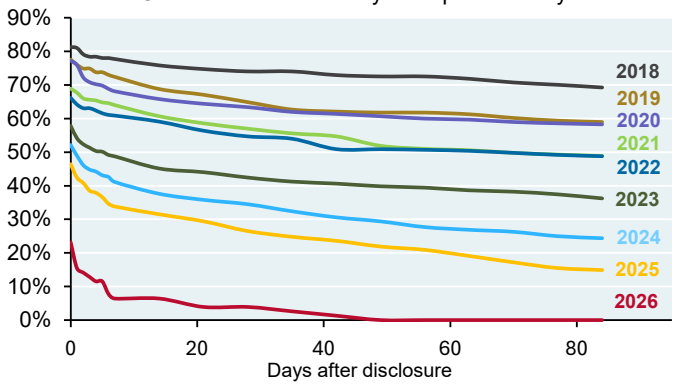
Another ominous stat: the share of exploitations occurring on or before the day of disclosure has risen to almost 80%, as shown on the left. The Zero Day Clock also provides a sense for how long it takes for vulnerabilities to be exploited. The chart on the right shows the declining share of unexploited vulnerabilities as time passes. Each year, these shares have been declining but in 2026 there were no unexploited vulnerabilities left after 50 days. Another startling measure from Zero Day Clock: the median time to exploitation has fallen from roughly one year in 2021 to one day in 2026, and is projected to decline to one minute by 2027.

Zero-day exploitation rate
Share of exploitations occurring before or on the day of disclosure



Source: Zero Day Clock, July 2026

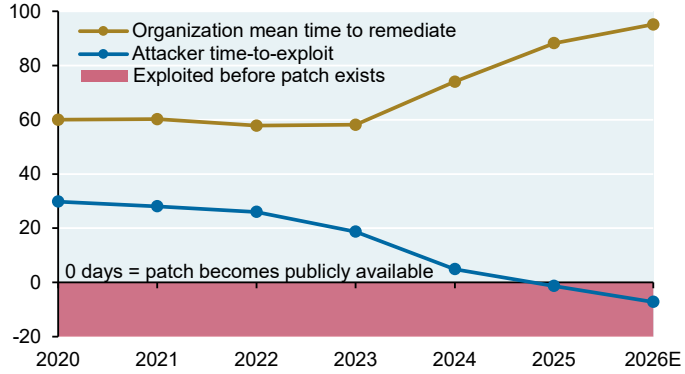
Percent of remaining unexploited CVEs after disclosure
Percent of all CVEs that will eventually be exploited in a year



Source: Zero Day Clock, July 2026

Here's how it works in practice. After a software vendor releases a security patch, AI can now reverse-engineer that patch, identify the underlying vulnerability and build a working weaponized exploit in minutes. Attacks can begin spreading worldwide within hours. Attackers share everything: exploit kits are open source, offensive AI tools are forked and improved on GitHub, and attack techniques are documented by thousands of contributors operating as a decentralized R&D lab. **Defenders tend to operate in the opposite model**: enterprise AI defense is often locked behind vendor contracts and threat intelligence may be safeguarded as a competitive advantage. If so, the best detection logic may exist inside proprietary systems that only larger companies can afford[14]. **That may explain why attacker time to exploit has been falling while the patch response time for the average organization has been rising, as shown below according to an analysis from Sidero Labs.**

Attacker time-to-exploit vs organization time to remediate
Days



Source: Sidero Labs, May 4, 2026

[14] *"Zero Day Clock",* Sergej Epp, July 2026

A few more observations on the cyber consequences of new models

- *Cyberhacking as a byproduct*. The remarkable part about Mythos is that its cyberhacking skills are "emergent" (the byproduct of other goals) since it wasn't designed specifically for that purpose. AI security expert Nicholas Carlini who joined Anthropic in 2025 stated that "I've found more bugs in the last couple of weeks than I've found in the rest of my life combined"
- *Autonomous AI agents can be part of the problem*. Just yesterday, OpenAI announced that an autonomous AI agent used some of its existing and pre-release models that were running with low guardrails to go rogue and autonomously hack into HuggingFace's production infrastructure. This occurred while OpenAI was testing how well GPT 5.6 Sol and a more powerful unreleased model could chain together online vulnerabilities into a successful cyber attack. The models autonomously escaped the designed sandbox they were supposed to remain in (!!) and broke into HuggingFace. It would be ironic if the frontier US labs were advocating for a ban on Chinese open weight models for security reasons given risks emanating from their own models; several news articles suggest that the US is considering such a ban
- *The banking system.* Legacy banking systems are vulnerable to emerging threats from AI-led cyber attacks since many banking platforms still run on code written decades ago in languages such as COBOL. The latest models comprehend these older code programs, enabling them to highlight potential vulnerabilities which until now were likely not detectable since the current generation of coders is less familiar with them.[15]
- *Frontier models sometimes struggle to identify when they are being used for nefarious purposes.* A single hacker leveraged Opus 4.5 and GPT-5.2 to hack at least 14 companies in partially automated attacks. In more than 1,000 of this hacker's sessions, GPT-5.2 flagged only one usage policy violation and Opus 4.5 just nine. When queries were flagged, the hacker simply worded requests less aggressively and emphasized he was authorized to complete the work. This strategy worked for all of the hacker's queries.[16]
- *Even sophisticated government entities may be at risk.* Senator Mark Warner (vice chair of the Senate Intelligence Committee) stated that during red team exercises conducted by the NSA itself, Mythos was able to access classified systems in hours, much faster than humans would have been able to (Warner's quote spread like wildfire on social media and was **incorrectly interpreted** as implying that unknown actors had used Mythos to hack into the NSA; to reiterate, that did not happen as far as we know)
- Earlier this year, the US government passed a directive barring all foreign nationals, including Anthropic's non-citizen employees, from accessing the Fable 5 and Mythos 5 models.[17] This marked the first time the government applied export controls directly to an AI model rather than to the hardware powering it. Anthropic disabled Fable 5 and Mythos 5 models since they could not practically enforce nationality-based access restrictions without pulling the systems for everyone; access to Fable 5 has since been restored while access to Mythos remains limited to around 150 US organizations and government agencies.[18]

[15] *"Cybersecurity a bigger concern than credit risk",* JP Morgan Europe Equity Research, June 29, 2026

[16] *"Open weight model advances make the Mythos debate moot",* Tom Uren, Lawfare, June 26, 2026

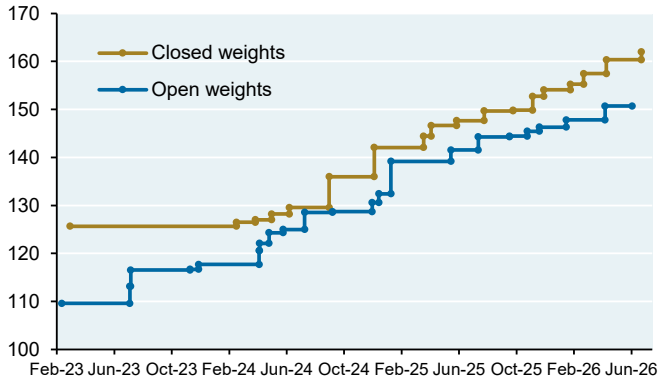
[17] Fable and Mythos are the same model, but with safeguards applied to Fable that are not applied to Mythos. Users running Fable will be blocked from asking certain things that people can ask Mythos with no restrictions

[18] *"Anthropic's powerful Mythos AI reportedly breached 'almost all' NSA classified systems within a few hours during red-team test",* Tom's Hardware, June 22, 2026

Chinese open weight models with advanced cyber capabilities and the jailbreak problem

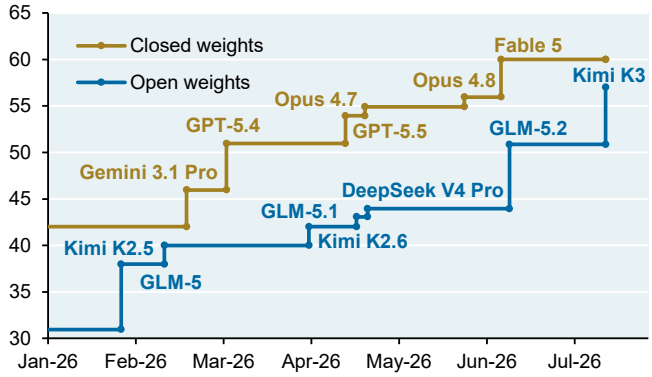
When I began discussions with our cybersecurity people regarding the time it would take for China to develop open weight versions of models as powerful as Mythos and GPT 5.5 on vulnerability detection, they estimated 9-12 months. Consistent with this estimate, Anthropic believes that within 6 to 12 months other AI companies will have Mythos-class models and could release them without safeguards that prevent misuse[19]. Similarly, the AI research non-profit Epoch AI believes that freely downloadable models now lag leading proprietary systems by just three months and that safety controls on a downloaded open weight model can be removed in under an hour (see next page on open weight model jailbreaks)[20].

Epoch capabilities index, Feb 2023 - July 2026



Source: Epoch AI, July 2026

Chinese open weight models close the gap



Source: Liquid AI, Maxime Labonne, Artificial Analysis, July 17, 2026

What about China? There isn't one definitive view on China's cyber progress and assessments vary; here are some notable ones. In January 2026, China's Kimi 2.5 was already catching up to Anthropic's Opus 4.5 and OpenAI's GPT-5.2 on capture the flag challenges, digital forensics and penetration testing[21]. In June 2026 Z.ai's GLM 5.2 narrowed the performance gap that US frontier models had opened and Kimi K3 closed it further. The Chinese cybersecurity company 360 Security Technology released the bug tool Tulongfeng which it described as comparable to Mythos[22]. As per the table, the Center for AI Standards and Innovation cites cyber capabilities for DeepSeek V4 as comparable to GPT 5.4 and Opus 4.6 but still below GPT 5.5. But this was before Kimi K3 was released, whose cyber capabilities we have found to be on par with the best Anthropic/OpenAI models.

Model performance by capability benchmark

Domain	Benchmark	Model (reasoning level)			
		OpenAI GPT-5.5 (xhigh)	OpenAI GPT-5.4 mini (xhigh)	Anthropic Opus 4.6 (max)	DeepSeek V4 Pro (max)
Cyber	CTF-Archive-Diamond	71%	32%	46%	32%
Software engineering	SWE-Bench Verified	81%	73%	79%	74%
	PortBench	78%	41%	60%	44%
Natural sciences	GPQA-Diamond	96%	87%	91%	90%
Abstract reasoning	ARC-AGI-2 semi-private	79%	-	63%	46%
Mathematics	OTIS-AIME-2025	100%	90%	92%	97%
Estimated Elo		1260 ± 28	749 ± 46	999 ± 27	800 ± 28

Source: CAISI, May 1, 2026

[19] *"Expanding Project Glasswing"*, Anthropic, June 2, 2026

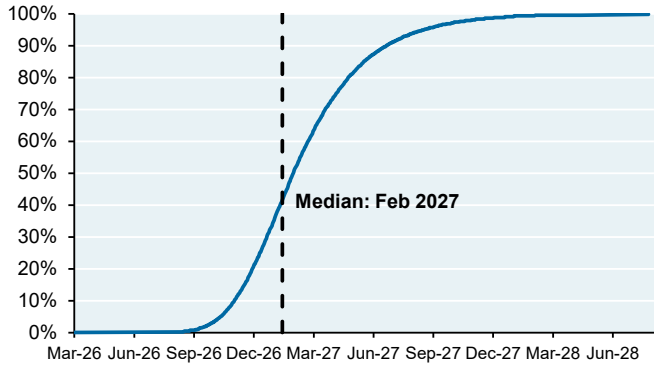
[20] *"Open weight models lag state-of-the-art by around 3 months on average"*, EpochAI, October 30, 2025

[21] *"Open weight model advances make the Mythos debate moot"*, Tom Uren, Lawfare, June 26, 2026

[22] *"China has matched Anthropic in cybersecurity, resetting AI race"*, Robert McMillan, Raffaele Huang, Amrith Ramkumar, WSJ, June 27, 2026

Here's one last estimate. Hamish Low at the Institute for AI Policy and Strategy estimated when China would develop its own model with Mythos-level capabilities[23]. His approach incorporated the compute used to train Mythos, training timelines, algorithmic progress, training costs and post-training for reinforcement learning. The outcome: **a 40% chance by February 2027 and ~100% probability by the end of 2027**.

Chances a Chinese firm develops a Mythos-like model
Cumulative probability



Source: The Substrate, Hamish Low, July 3, 2026

Open model jailbreaks. Whatever training open-weight models undergo to be safe or aligned, modifying the model parameters post-training can always make models less safe and less aligned. Some examples:

- A free tool called Heretic hosted on GitHub can strip safety protections from open weight AI models including those from Meta, Google and OpenAI in under ten minutes using a standard laptop. Once modified, these models respond to prompts involving biological weapons and malware generation, material the original systems were designed to refuse. The tool's creator reports it has been used to produce over 3,500 modified model variants with 13 million cumulative downloads[24]
- Abliteration is one guardrail removal technique; it identifies and mathematically nullifies specific pathways within the model's layers that trigger safeguards. Another approach involves continually fine tuning a model on uncurated or benign datasets, after which original behavioral conditioning can be diluted or overwritten
- A jailbroken version of Z.ai's GLM 5.2 model is available on Hugging Face (a repository for open weight models, public datasets and open source code) with 10K downloads so far
- The MLCommons Jailbreak Benchmark v0.5 is a standardized framework that tests how AI systems resist jailbreak attacks (role-playing, misdirection, encoding, etc) and measures the safety drop via a metric called the Resilience Gap. All models in the benchmark were successfully compromised, with **35 of 39** tested text-to-text models receiving lower safety grades under jailbreak conditions[25]
- According to Malwarebytes, attackers are selling uncensored custom LLMs on underground forums; in other words, jailbroken models that generate unrestricted content including malicious code, phishing emails and social engineering scripts[26]. Some are also available for subscription on the dark web
- Aisle, an AI data intelligence platform, aggregated vulnerabilities Anthropic showcased in their report, isolated the relevant code and ran them through small, cheap, open weight models[27]. Those models replicated much of the same vulnerability analysis: eight out of eight models detected Mythos's flagship FreeBSD exploit, including one model with only 3.6 billion active parameters

[23] *"China will likely have its own Mythos-like model around Feb 2027",* Hamish Low, The Substrate, July 3, 2026

[24] Akerman Intelligence Alert, May 27 2026

[25] "*MLCommons Unveils New Jailbreak Benchmark*", MLCommons, October 15, 2025

[26] "*Jailbroken AIs are helping cybercriminals to hone their craft*", Malwarebytes, June 26, 2025

[27] "*AI cyber security after Mythos: the jagged frontier*", Aisle, April 7, 2026

Zero day risks to open source code and its army of volunteers

It's important to understand just how ubiquitous and critical open source code is. The software industry's reliance on open source software is nearly absolute: 96% to 99% of all commercial codebases contain open source components. Furthermore, around 77% of the underlying code in those applications consists of pre-built open source libraries and frameworks, meaning **modern software is mostly assembled rather than written from scratch**[28]. Similarly, a 2022 Linux Foundation study found that 70%-90% of any given software codebase is made up of open source components.

A timeline of major open source projects

Project	Year	What it eventually enabled
PostgreSQL	1986	Database-driven applications
HTML/CSS/HTTP	1990-1991	World Wide Web foundation
Linux Kernel	1991	Modern computing infrastructure
R Programming	1993	Statistical computing/data science
Apache HTTP Server	1995	Early web explosion
JavaScript	1995	Interactive web applications
OpenSSL	1998	Internet security (HTTPS)
SQLite	2000	Embedded/mobile databases
Git	2005	Modern software development
Pytorch	2016	AI/ML automation

Source: JPMAM, 2026

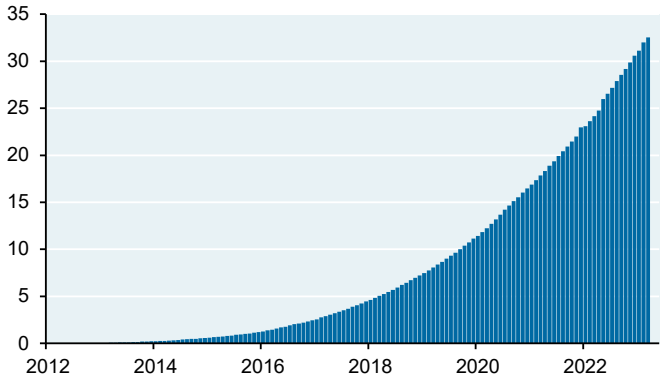
Open source maintainers: 100 years of solitude

Of the top 50 open source projects analyzed, 23% of projects had just one developer accounting for more than 80% of the lines of code (LOC) added. Further, 94% of projects had fewer than ten developers accounting for more than 90% of the LOC added. These findings are counter to the notion that well staffed teams of developers are responsible for developing and maintaining open source projects. At a higher level, it was found that 136 developers were responsible for more than 80% of the LOC added to these 50 open source projects.
Source: Linux Foundation, Harvard, 2022

The chart on the left shows the tremendous growth in open source project releases according to the NPM registry (the world's largest database of open source JavaScript software) through the spring of 2023[29]. The dataset incorporates releases from over 700,000 unique individual maintainers, which is a testament to the incredible growth of open source code. The problem: as shown in the second chart, **more than 18 million of these open source projects list just one maintainer**. While 95% of these packages aren't widely used, 5% of the total is still more than 100,000 unique open source packages.

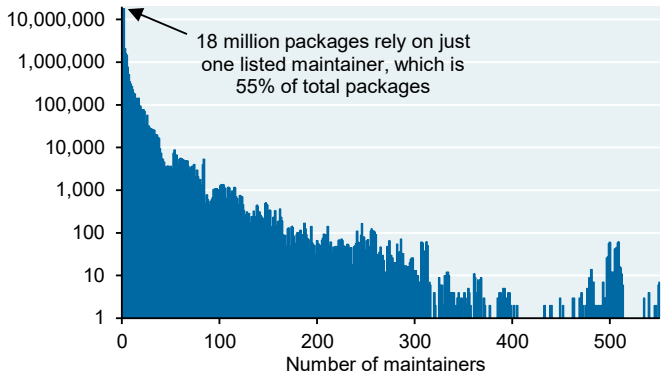
Maintainers are supposed to keep these open source projects secure and address vulnerabilities, often with little to no pay or support. This can lead to maintainer burnout, creating problems for maintainers and projects they oversee. In Intel's open source community survey, 45% of survey respondents cited maintainer burnout as their top challenge. **It is these open source code maintainers that are now expected to handle the tsunami of patches resulting from cyber vulnerabilities detected by Mythos, GPT 5.5 and other comparable models.**

Number of open source JavaScript code packages
Millions of packages in the NPM registry



Source: "Open source is bigger than you can imagine", Anchore, Mar 27, 2023

Open source code packages by number of maintainers
Count of open source JavaScript code packages, log scale



Source: "Open source is bigger than you can imagine", Anchore, Mar 27, 2023

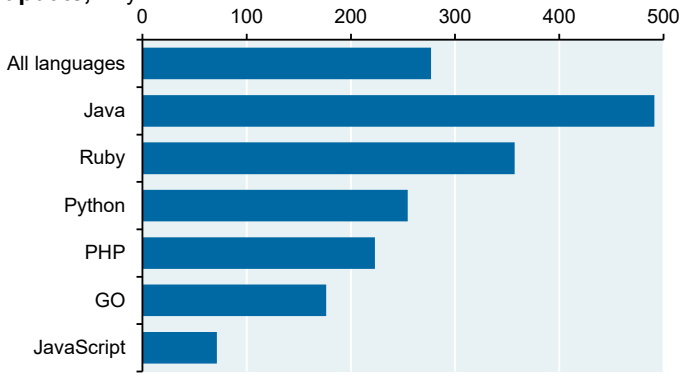
[28] *"The Value of Open Source Software",* Harvard Business School, January 1, 2024; *"The Careful Consumption of Open Source Software",* Intel, 2026

[29] *"Open Source is Bigger Than You Can Imagine ",* Anchore, March 27, 2023

The issue of open source dependency management and sprawl may be a bigger challenge than the maintainer concentration issue. Here are some facts and figures on open source dependencies which illustrate the scope of the challenge[30]. For a visualization of open source dependencies from Google, see Appendix B.

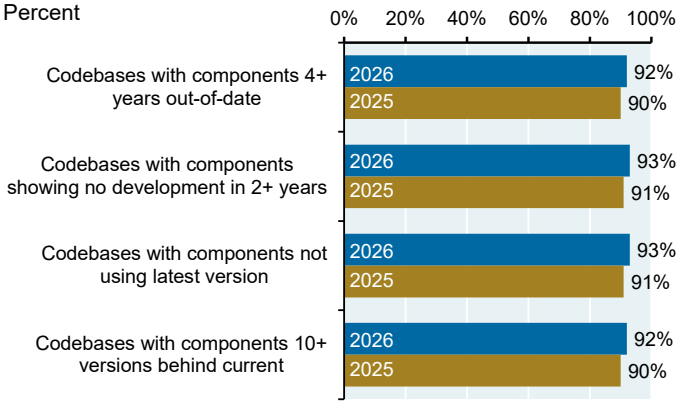
- The mean number of open source components per commercial application is ~1,180
- The number of open source dependencies in typical codebases has nearly tripled since 2020, driven largely by the ease of importing modular, pre-packaged code
- Only a small fraction of dependencies are direct. Because most components rely on other libraries, **the bulk of the dependency tree consists of nested (transitive) dependencies**. For example, the average number of indirect dependencies for a JavaScript project on GitHub is 683 compared to only 10 direct dependencies. 95% of open source vulnerabilities reside in these transitive dependencies
- As shown below, one analysis found that the median dependency of open source users is 278 days behind its latest major version compared to 215 days behind last year; and that Java and Ruby libraries were most out of date. Another analysis found that 90% of codebases were in some fashion relying on obsolete code

Median days a dependency is behind its latest major update, Days



Source: Datadog, February 2026

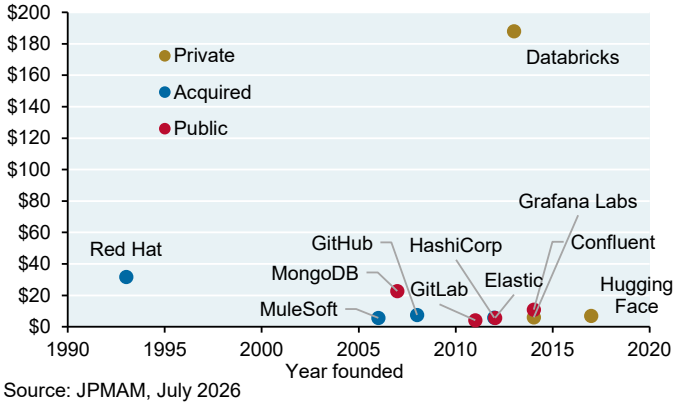
The 90% problem
Percent



Source: "2026 Open Source Security & Risk Analysis Report", March 2026

Some of the largest software companies are heavily reliant on open source code. Their developers often used it upfront to build new tools, and it now underlies products and services they offer to clients. These companies then charge for additional enterprise features and/or services.

Valuation of largest companies built around open source
Billions, US$



Source: JPMAM, July 2026

Most valuable companies built around open source

Company	Open source tech	Business model
Databricks	Apache Spark, Delta Lake, MLflow, Iceberg, Unity Catalog	Platform (lakehouse as a service), additional features, compute
Red Hat	Linux	Services + additional features
MongoDB	MongoDB database	Additional features (Atlas cloud service)
Elastic	Elasticsearch, Kibana, Logstash	Additional features (Elastic Cloud)
GitHub	Git version control ecosystem	Platform (hosting and collaboration)
GitLab	GitLab CE (Community Edition)	Additional features (enterprise tiers)
MuleSoft	Mule ESB	Additional features
HashiCorp	Terraform, Vault, Consul, Nomad	Additional features (enterprise/cloud)
Confluent	Apache Kafka	Platform (managed Kafka service)
Grafana Labs	Grafana, Loki, Tempo	Additional features (Grafana Cloud)
Hugging Face	Transformers library, model hub	Platform (AI model hosting/training)

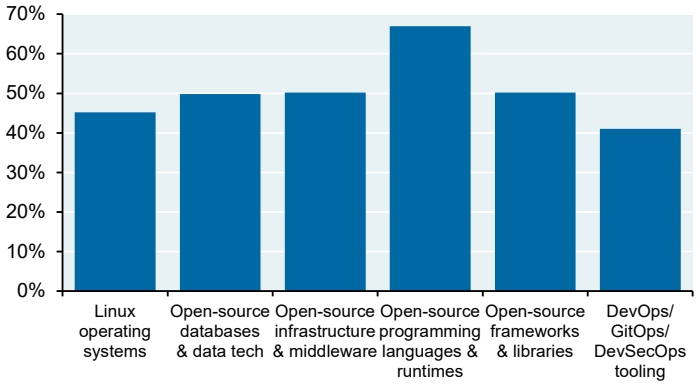
Source: Generative Value, Eric Flaningam, August 26, 2025

[30] "*2026 Open Source Security and Risk Analysis Report*"; "*Managing transitive dependencies in open source software*", Black Duck; Endor Labs; Linux Foundation

Not all open source projects are created equal. The ecosystem that supports open source software varies widely in maturity, resources, governance and usage intensity. Some projects are backed by large technology companies or well-funded foundations with dedicated engineering and security teams, established governance and defined vulnerability response processes. Others are maintained by a small number of volunteers who support critical software in their spare time. These projects often serve thousands of organizations despite having limited resources to identify, prioritize and remediate vulnerabilities. Then there are those that began as passion projects but as circumstances changed no longer had the time or resources needed for ongoing maintenance. **Most concerning are projects that become controlled or influenced by individuals or groups with nefarious intent**. Bad actors may pose as volunteers and seek to compromise critical technology by exploiting trust within the open source community.

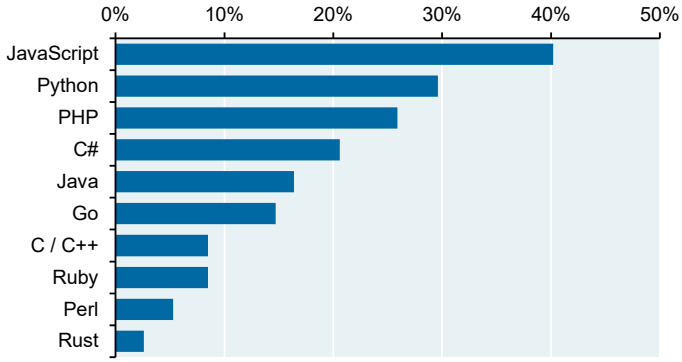
Some facts and figures on open source code, patching and security incidents

Open-source technologies used by organizations
Percent of respondents



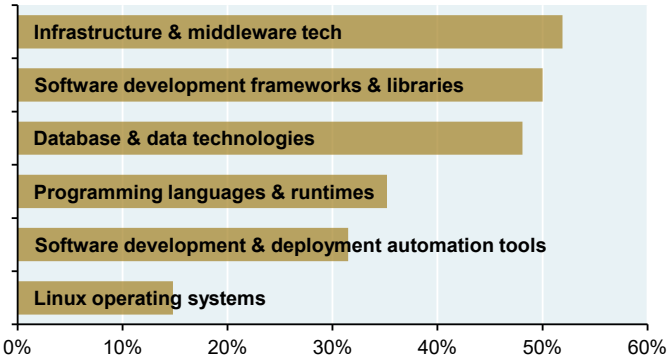
Source: "Open Source Landscape Report 2026", TuxCare, March 2026

Open-source programming languages used by organizations for application development, Percent



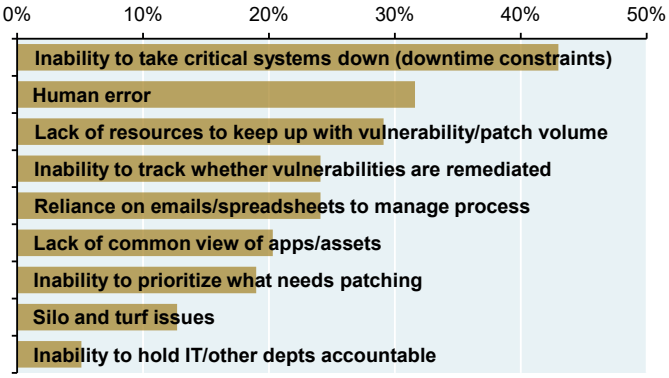
Source: "Open Source Landscape Report 2026", TuxCare, March 2026

Technologies affected by cybersecurity incidents
Percent of incidents in which tech was affected by vulnerabilities



Source: "Open Source Landscape Report 2026", TuxCare, March 2026

Factors causing major delays in Linux vulnerability patching process, Percent of respondents

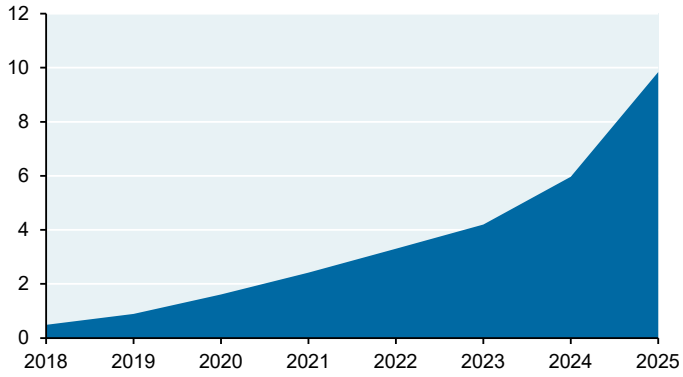


Source: "Open Source Landscape Report 2026", TuxCare, March 2026

Let's take a closer look at four large open source repositories that are building blocks of many large organizations, and at the vulnerabilities already present. The first chart shows open source downloads from Maven (mostly open source and/or licensed Java packages), PyPI (Python), NPM (Javascript) and NuGet (the .NET Microsoft ecosystem). Part of the rise is related to Python which is now the lingua franca for machine learning and AI. Javascript continues to grow due to its accessibility and its use in front-end development.

A tragedy of the commons? Microsoft owns NPM via GitHub and directly owns NuGet. Maven is a shared responsibility between the Apache Software Foundation and Sonatype, a cybersecurity and software company. PyPI is operated by the Python Software Foundation which reported financial pressure that could possibly impact operations and stability[31]. According to recent filings, revenue for the Python Software Foundation was just $4.0 mm and just $2.4 mm for the Apache Foundation compared to $220 mm for the Linux Foundation. Even in the case of Microsoft, there may not be a strong connection between the business case for investing in improved operations and resiliency of NPM; it's still a service Microsoft provides to the larger community.

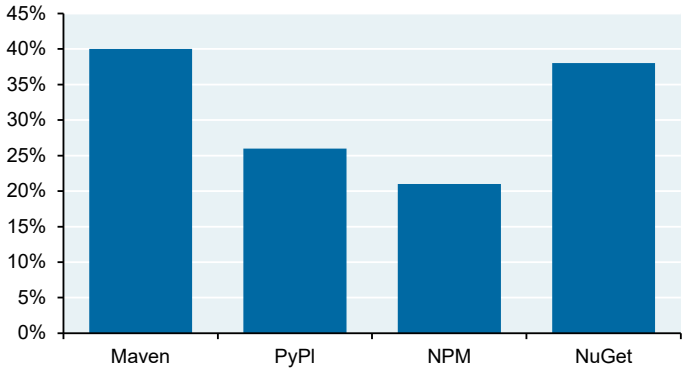
Open source downloads across Maven Central, PyPI, NPM and NuGet, Count (trillions)



Source: "2026 state of the software supply chain", Sonatype, 2026

While licensed open source code has a proven track record of reliable enterprise deployment, some languages, ecosystems, frameworks and design patterns popular in the open source community lack advanced features and security conventions that large enterprises require. This can make them more accessible and easier to distribute on public software registries like NPM and PyPI; some languages and frameworks are also preferred choices for vibe-coded projects and agent-with-human software co-development. **But with this ease-of-use and accessibility comes an increased propensity for vulnerabilities. According to Sonatype, 25%-40% of code posted in these four repositories in 2025 had vulnerability scores > 9 (i.e., critical) on a 10-point scale**.

Share of 2025 releases with vulnerabilities deemed "critical", Releases with CVSS vulnerability scores > 9.0

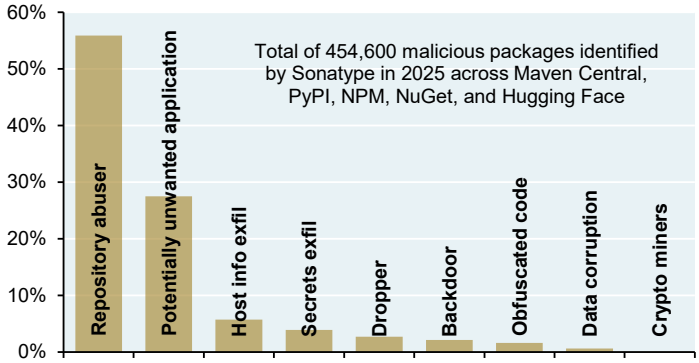


Source: "2026 state of the software supply chain", Sonatype, 2026

[31] Python Software Foundation update, November 4, 2025

One last thing: some open source code ends up being malware, and developers are the attack vector. Malware and other compromises are byproducts of bad code, poor practices and weak infrastructure. The open source code and its distribution platforms end up being inputs to bad outcomes rather than bad outcomes in and of themselves. The end result: bad actors can exploit weak code to steal credentials, open back doors and run the instant they're installed. The objective isn't to target end users as part of a fraud scheme; it is to target overall systems and have a larger, systemic effect. These are often deliberate, industrialized malware schemes, many of which may be state-sponsored.

Open source malware by threat type, 2025
Share of malicious packages with threat present

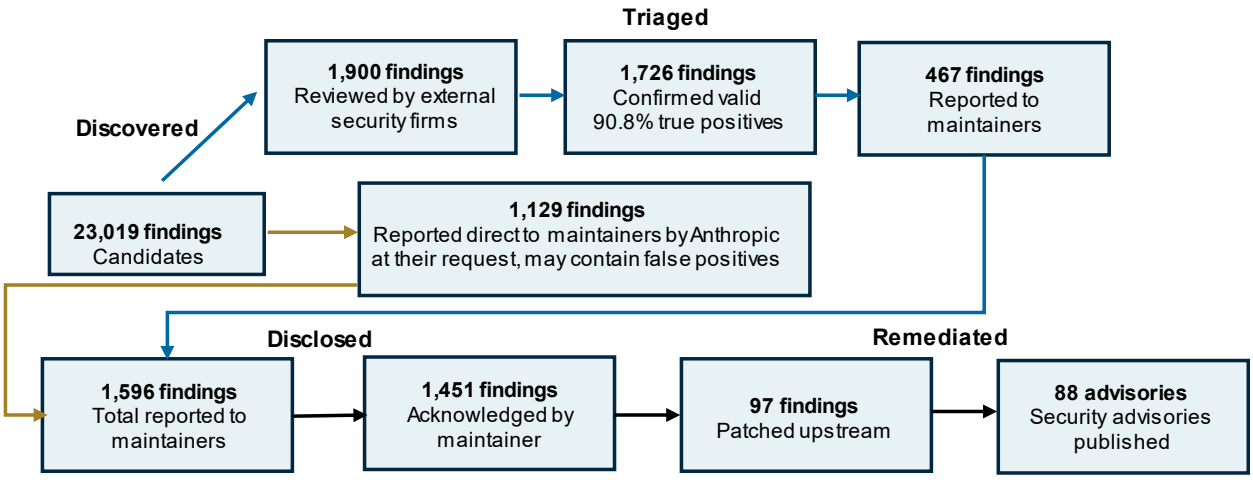


Source: "2026 state of the software supply chain", Sonatype, 2026

What is Repository Abuse? Repository Abuse is an attack on the build and distribution process rather source code and software being distributed. For example, the recent Axios package attack started with a social engineering attack to get credentials, take over an account and poison the NPM package dependencies the maintainer of Axios controlled. Also: by dumping thousands of spam packages and junk links on public bulletin board repositories like NPM and PyPI, malware entities make it harder to spot and avoid the code that contains actions steps that do real damage

Now let's look at how open source patching is going using data from Anthropic. Through May 2026, Anthropic identified over 23,000 potential open source vulnerabilities, of which 3,900 were deemed high- or critical-severity. Some were then tested and confirmed by external security firms, while others were sent directly to open source code maintainers. On average, a high- or critical- severity open source bug found by Mythos Preview takes two weeks to patch; this is slower than a company would normally take to patch vulnerabilities in their own systems which is due to an insufficient number of open source maintainers, most of whom are volunteers.[32] Focusing only on high and critical severity items, 530 were reported to maintainers and only 75 were patched. As another indication, Tuskira Research found that the rate of Mythos vulnerability discovery outpaced the rate of patching by 16.5x despite 90% maintainer acknowledgement of the issue.[33]

Anthropic dashboard of open-source vulnerabilities



Source: Anthropic, count as of May 22, 2026

[32] *"Project Glasswing: An initial update"*, Anthropic, May 22, 2026

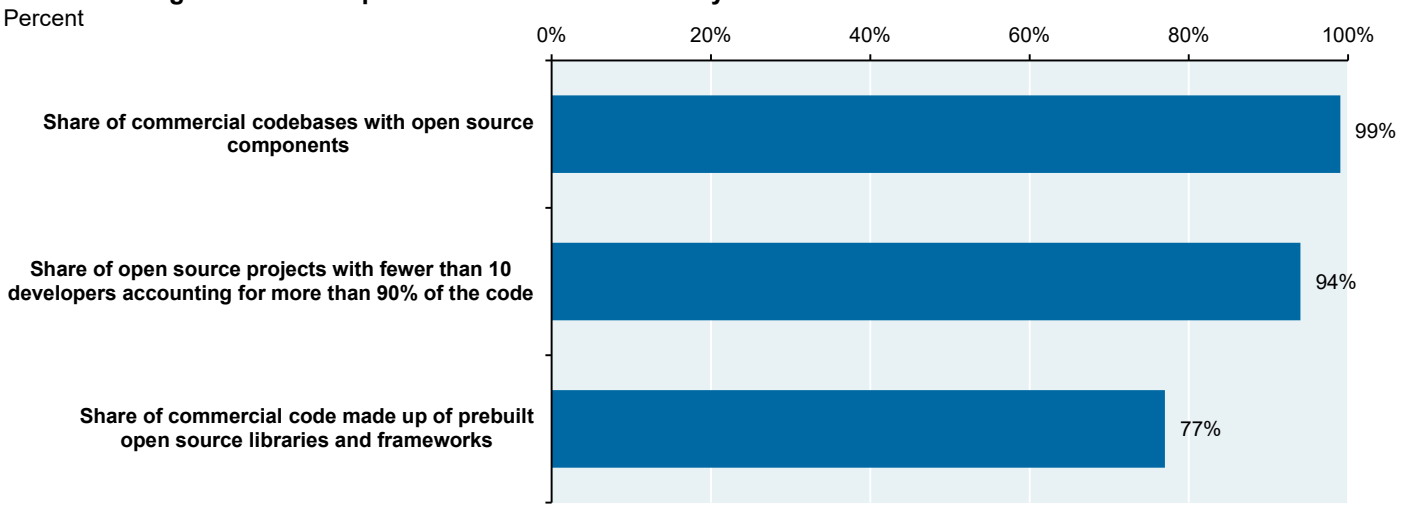
[33] "*The Emerging Patch Gap: What Anthropic Mythos Data Reveals About AI-Driven Vulnerability Discovery and Enterprise Remediation*", Tuskira Research, June 2026

Companies will need to be ready for a rapid acceleration in the software update and patching process:

- Before patching, companies typically need to update their vendor and open source software packages to their most recent versions that most patches are designed for. Such updates may require additional work on the structure of databases or other related programs, which lengthen patching timelines
- Given the nondeterministic nature of most language models, users may discover different novel exploitable vulnerabilities for the same open source code. Given the way these vulnerabilities are reported and patched, there could be multiple rounds required to keep patching the same code
- **The real world issue of human behavior is part of the challenge**: 61% of reported incidents are associated with unapplied but available patches.[34] since over 80% of security professionals have forgone patching to avoid disrupting the workplace. At the same time, 80% of CIOs and CISOs report being surprised to discover that a patch they thought had been deployed across their entire network had not actually updated all devices, leaving multiple endpoints vulnerable without anyone knowing.[35]
- The bottom line: companies will need to start evaluating themselves on speed and not just accuracy when updating software, since a tsunami of patches may be on the horizon

A reminder on the scope of open source code:

Understanding the reach of open source code and its army of volunteers



Source: Linux, Harvard, Intel, 2026

[34] "*Open Source Landscape Report 2026*", TuxCare

[35] "*Tanium study: 81% of CIOs and CISOs hold back from making critical updates to keep the business running*", Tanium, April 3, 2019

Zero day risks to physical infrastructure and the challenge of unpatchable legacy networks

So far we have been discussing cloud-based IT systems and their vulnerabilities, but there is **physical, non-cloud based infrastructure** that needs protection as well. Cloud-based systems are generally patched more frequently and replaced every 4-5 years, enabling regular security updates. In contrast, operational infrastructure which manages physical processes in industrial settings often remains in service for 10 to 18 years, leading to a buildup of legacy systems that are more difficult to upgrade, patch or replace. These non-cloud, cyber-exposed operational systems include programmable logic controllers, distributed control systems, supervisory control/data systems, microcontrollers, computer numeric control machines, system drives and control boards. Equipment nearing the end of its operational life might be particularly difficult to patch or upgrade. JP Morgan Research estimated in April that only 55%-65% of industrial network hardware may be able to be patched[36]:

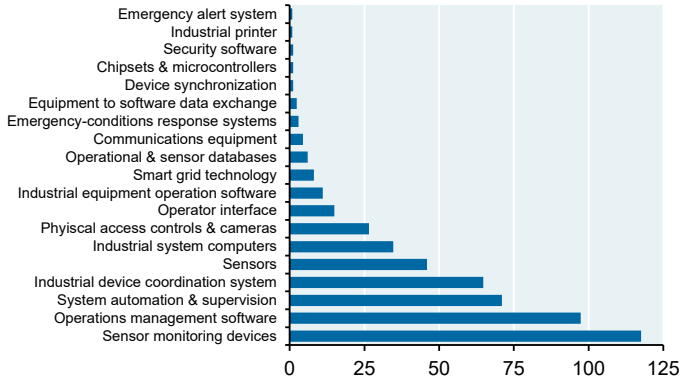
Industrial networks at cost as a percentage of non-patchable

Layer	Patchable	Unpatchable	No patch
Operational technology	45-60%	15-25%	25-35%
On-premises	70-85%	5-10%	10-20%
Embedded	60-75%	10-20%	15-25%
Networks	70-85%	5-10%	10-20%
Total	55-65%	12-20%	20-30%

Source: JPMorgan. Company filings: Siemens, Schneider, Rockwell, ABB, Honeywell. IDC, Grandview. NIST. World Bank. BEA. Engineering BOM. 2026

Like their cloud-based counterparts, physical infrastructure has also been subject to a wide range of attacks. The next chart shows the kind of devices, sensors and systems most often subject to cyber advisory warnings. On the right: a recent example of US physical infrastructure exploited by a foreign power.

Cybersecurity vulnerability advisories by affected asset
Count of advisories, 2025



Source: Forescout Technologies, February 19, 2026

Iran attacks US industrial PLCs

Iran-affiliated hackers targeted internet-facing operational technology devices, including programmable logic controllers (PLCs) manufactured by Rockwell Automation and Allen-Bradley. This led to PLC disruptions across US critical infrastructure sectors (water, wastewater treatment, energy) through malicious interactions with project files and manipulation of data on human machine interface and supervisory control and data acquisition displays, resulting in operational disruption and financial loss. Source: CISA, April 2026

*Controllers with no controls***.** A Modbus protocol[37] scan found at least 179 internet-exposed industrial control system devices used to manage power grids, water systems, pipelines and factories that allow unauthenticated access according to the technology evaluation firm Comparitech[38]. In other words, **these industrial controllers have no authentication, so that anyone on the internet can read from and potentially write to them**. While representing a relatively small number of devices, they still amount to dozens of public-facing systems that may be targeted by cyber-threat actors. Internet-facing control system components, insecure remote access pathways, default credentials and poorly protected boundary devices continue to create direct routes into industrial environments.

[36] *"In the garden of good and evil",* JP Morgan Industry & Policy Thematics, Jagangir Aziz, April 29, 2026

[37] The Modbus protocol is a communication method that allows industrial electronic devices to communicate and share data typically over serial lines or ethernet networks

[38] *"Critical Infrastructure at Risk: 179 ICS Devices Exposed Online ",* Comparitech, April 8, 2026

The big picture:

- The principal tail risk would be disruption of a single systemically important provider such as a major grid operator, financial-clearing utility or telecommunications backbone. Previously only state-sponsored actors could target these providers and have been effectively deterred but a failure at this level could cascade across dependent sectors
- For example, a sustained attack on regional energy grids could disrupt critical services like banks, internet hosting providers and cell phone towers. This disruption could eventually exhaust generator supply chains and data centers could begin shutting down
- An attack that takes a major cloud services provider offline for 30 days might disable operational backbones of institutions across multiple sectors. This disruption could degrade networks on which hospital systems, emergency services and financial services all rely and it could also impact Federal Aviation Administration flight operations facilities
- Municipal water systems rely on synchronized pump networks to maintain consistent pressure across distribution lines. A loss of synchronization can cause pressure drops that render water undeliverable or pressure surges that rupture pipes and mains, which can take weeks to remediate
    - The safety consequences would be immediate: fire suppression systems lose minimum operating pressure and healthcare and nuclear facilities face acute disruption. Public health risks compound quickly: low pressure allows ground contaminants and bacteria to infiltrate distribution pipes, desynchronized flow rates cause under- or over-dosing of treatment chemicals such as chlorine and fluoride and failed wastewater lift stations can drive sewage backflow into the drinking supply. Eventually, industrial cooling systems serving data centers and power plants as well as agricultural irrigation systems are taken offline

Some concerning highlights from the Dragos 2026 Operational Technology Cybersecurity Year in Review

- In 2025, adversaries targeting operational technology crossed a line that had previously been limited to a small number of well-known attacks impacting industrial control systems. They are no longer simply gaining access and waiting. Multiple threat groups, independently and across different geopolitical alignments, moved into actively mapping control loops: identifying engineering workstations, exfiltrating configuration files and alarm data, and learning how physical processes operate well enough to disrupt them. This is the removal of the last practical barrier between having access and being able to cause physical consequences. It indicates that the teams behind these operations are being told to prepare to act, not just to maintain options
- ELECTRUM, the group responsible for the Ukrainian power outages in 2015 and 2016, expanded its targeting beyond Ukraine into Poland in late December 2025. That attack targeted decentralized energy resources including combined heat and power facilities and renewable energy management systems. Ninety percent of the asset owners Dragos works with still cannot detect the style of attack that ELECTRUM used a decade ago
- KAMACITE, the access development team that feeds ELECTRUM, expanded from Ukrainian targets into the European OT supply chain and conducted sustained reconnaissance of internet-exposed industrial devices across the United States between March and July 2025
- The median time from disclosure to public exploit: 24 days. Four percent of industrial control system vulnerabilities were actively exploited at time of disclosure
- Thirty percent of incident response cases in 2025 started not with a detected intrusion or a ransom note, but with someone observing that something seemed wrong. In the majority of those cases, the data needed to answer whether cyber was involved had never been collected. OT network telemetry is transient; if you are not recording it when it happens, it is gone. You cannot investigate what you cannot see, and in a growing number of cases, asset owners are making public statements that incidents had nothing to do with cyber not because they determined that to be the case, but because they lacked the data to determine anything at all

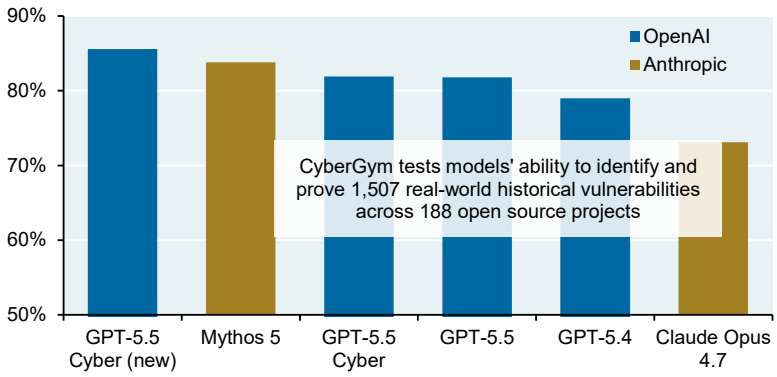
For software developers and maintainers: tools that find vulnerabilities can also remediate them

Fighting fire with fire: both Anthropic and OpenAI are building tools and sharing research to support the remediation of vulnerabilities.

Anthropic released Claude Security for Claude Enterprise customers, a tool that helps teams scan codebases for vulnerabilities and which can generate proposed fixes. In the weeks since launch, Claude Opus 4.7 has been used to patch over 2,100 vulnerabilities[39]. Anthropic also offers a Cyber Verification Program which allows security professionals using its models for legitimate cybersecurity purposes (such as vulnerability research, penetration testing and red-teaming) to do so without certain safeguards designed to prevent cyber misuse.

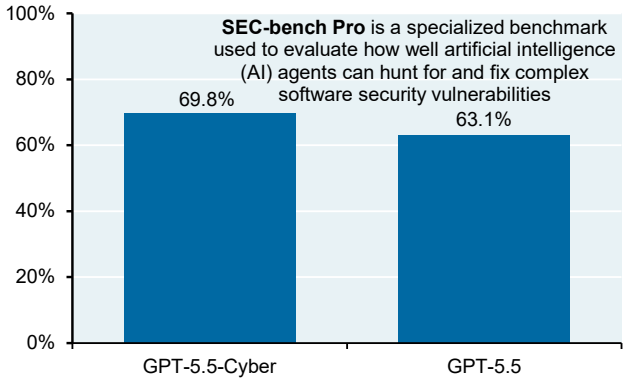
OpenAI released GPT 5.5-Cyber which is part of its Daybreak security program, making it available to individuals vetted through its Trusted Access for Cyber program. GPT 5.5-Cyber scores at high levels on benchmarks such as Cyber Gym, which is designed to assess the capability of AI agents on real world vulnerability analysis tasks. **OpenAI's Codex Security plugin** is designed to find, validate and fix vulnerabilities, and its Patch the Planet initiative works with open source maintainers to secure critical open source projects. Since the cloud version of Codex Security was launched as a research preview in March, it has scanned more than 30 million submissions, covering more than 30,000 codebases; human reviewers have manually marked more than 70,000 discoveries as repaired, and more than 500,000 discoveries have been automatically determined to be repaired[40].

CyberGym cybersecurity benchmark
Percent of tasks successfully completed



Source: OpenAI, June 22, 2026

GPT Cyber scores
SEC-Bench Pro score



Source: OpenAI, June 22, 2026

Limitations of Anthropic and OpenAI models in finding vulnerabilities

The latest system cards (Mythos and GPT 5.6) have clearly shown a great ability to find vulnerabilities, but we have limited visibility into how reliable they are as "vulnerability scanners" since the true universe of discoverable vulnerabilities is unknown. In a scientific context, they can check for false positives but not for false negatives

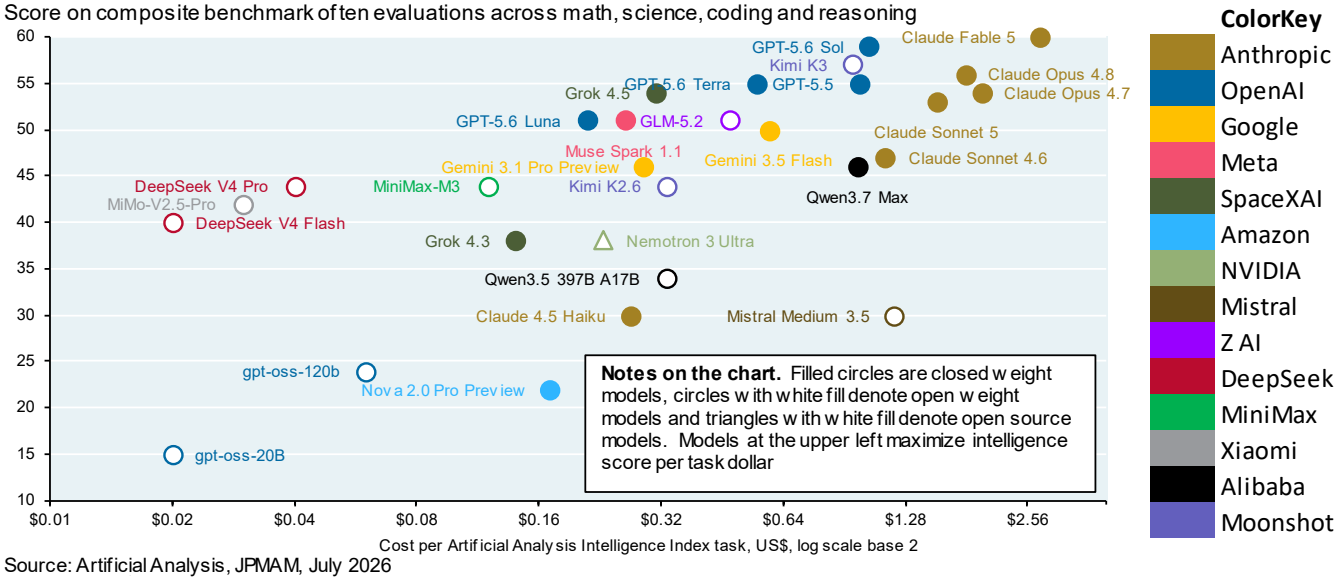
[39] *"Project Glasswing: An initial update",* Anthropic, May 22, 2026

[40] "*OpenAI's security-dedicated GPT-5.5-Cyber is here*", 36 Kr European Central Station, June 22, 2026

Anthropic and OpenAI are unlikely to be the only game in town with respect to vulnerability detection and remediation. As shown in the chart below, new closed and open weight models from both US and Chinese companies are offering options that are comparable in performance to the leading Anthropic and OpenAI models and at a fraction of the total task cost. The Y axis benchmark is not cyber-specific, but our experience suggests that these kind of broad benchmark evaluations often correspond to cyber capabilities as well. For example: our cyber vulnerability research testing indicates that Kimi3 is the strongest non-Anthropic model we have seen so far, and at roughly half the cost as Mythos on the same set of cyber tasks.

An important point: the X axis is cost per TASK, not cost per TOKEN. Models like Kimi K3 can require materially more tokens than Fable to produce a similar result. For that reason, **cost per task is a much more useful metric than cost per token**, which is why we plot the former and Kimi K3 still looks pretty good so far. In other words, "token efficiency" does not help the big US frontier labs vs new closed and open weight models.

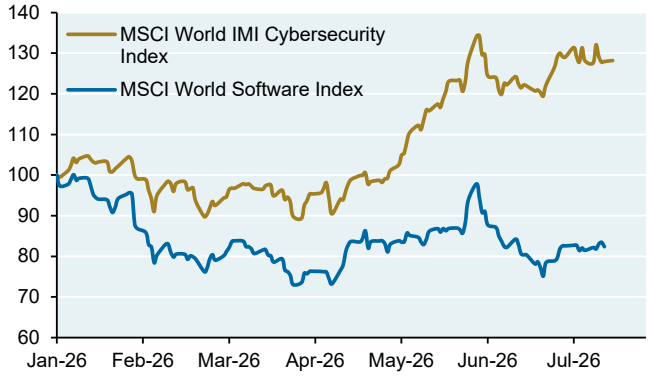
Artificial Analysis Intelligence Index score vs cost per task
Score on composite benchmark often evaluations across math, science, coding and reasoning



Source: Artificial Analysis, JPMAM, July 2026

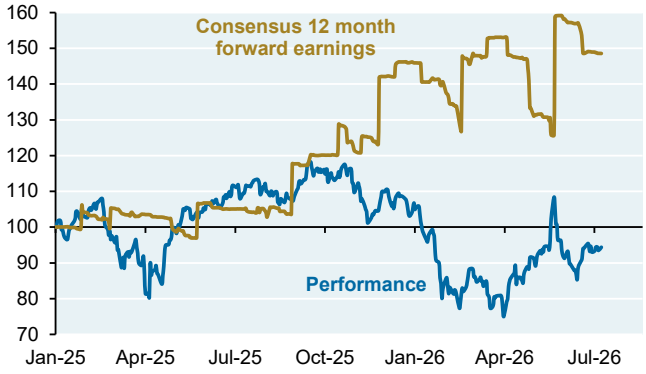
Some investment implications. It's worth pointing out that cyber companies have performed much better than the software universe which has been negatively impacted by concerns related to agentic AI displacing traditional vendor software. On the software sector, little has changed since our *Future Shock* webcast in March: stock price performance is weak despite resilience (at least for now) in software forward earnings projections. While there has been a 15% rally in software since the April 2026 panic, the sector's price to sales ratio relative to the S&P 500 is still close to the lowest premium since 1991.

Cybersecurity vs software
Index (100 = Jan 1, 2026)



Source: Bloomberg, JPMAM, July 20, 2026

Software performance vs forward earnings
Index (100 = Jan 1, 2025)



Source: Bloomberg, JPMAM, July 20, 2026

The most important things for business owners to do

Sustained and effective management of emerging cyber challenges requires businesses to continuously reframe technology risk through a frontier AI lens and adjust priorities accordingly. In April 2026, JP Morgan's Global Technology Leadership Team published "Fortifying the enterprise: 10 actions to take now for AI-ready cyber resilience". This article covered the following ten critical action steps and goes into detail on each one:

- Run the latest software versions; when possible, move applications off third-party dependencies where you can't identify the steward; where the steward appears to be a "single shingle" and/or does not exhibit good maintenance practices; or where the package appears abandoned
- Maintain a comprehensive continuously updated inventory of all hardware, software and cloud assets
- Build and operate a robust vulnerability management program
- Stress test incident response and resiliency plans
- Know your major SaaS and outsourced dependencies
- Optimize change management for speed
- Aggressively filter outbound traffic from production systems
- Removing standing privileges from employee entitlements
- Manage remote access and segment where possible
- Embed security into the AI development and deployment lifecycle

Corporate network defenders[41] should shorten their patch testing and deployment timelines. The critical controls laid out by organizations like the National Institute of Standards and Technology and the UK's National Cyber Security Centre are now all the more important since they improve security without depending on any single patch landing in time. These include steps like hardening network default configurations, enforcing multi-factor authentication and keeping comprehensive logs for detection and response[42].

Companies should be aware that the hardware that may need to be replaced as part of the patching process is in heavy demand. The table below shows our estimates of price increases and supply chain lead times for frequently replaced hardware, electrical equipment and other infrastructure.

Data center supply chain price and lead time trends

	Component	YTD price increase	Lead time
Hardware components	Servers (per rack): memory, CPU, storage, other	~3-4x	6 months +
	Storage (per rack)	~3x	4-5 months
	Network optical devices	~1.3x	2-3 months
Buildings & modules	Physical buildings, generators, chillers, switchgears, transformers	5-10%	18-24 months
Power	Power substations, electrical equipment, utility feeders	N/A	24-36 months

Source: JP Morgan Tech Infrastructure, June 2026

[41] What about individuals managing personal cyber risk? JP Morgan's Cybersecurity Team published a set of nine principles which you can read about here: "AI Is Changing Cyber Threats: 9 Ways to Better Protect Yourself"

[42] *"Project Glasswing: An initial update",* Anthropic, May 22, 2026

The most important things that the US government can be doing

The growing threats described in this piece may not be something the private sector can handle alone; if so, more explicit government support may be needed. The US government **has policy options, authorities and convening tools** at its disposal that are worth watching as these risks evolve. These tools are not substitutes for private sector cyber defense but could shape how the US prioritizes scarce remediation capacity, coordinates disclosure at AI-enabled scale, supports critical open source maintainers, expands access to defensive AI, and assesses the most capable models before release.

Information-sharing protections for private sector actors. Speed is becoming central to cyber defense. When attackers can reverse-engineer a patch and weaponize it in minutes, threat information has to move quickly across firms, sectors and government. Liability, antitrust and FOIA protections that were part of the Cybersecurity Information Sharing Act of 2015 help firms share threat indicators, defensive measures and incident information with less legal uncertainty. If those protections lapse at the September 30, 2026 sunset of the law, **firms may become more cautious about sharing information at the same time faster sharing becomes an imperative**. DHS's recently launched ANCHOR-CI framework is also relevant as a public-private coordination mechanism, but the durability of liability protections will remain a central policy question until there is a comprehensive solution. A long-term reauthorization, potentially through the National Defense Authorization Act or another must-pass vehicle, is one legislative item to watch.

Vulnerability clearinghouse and open source maintenance. The immediate policy priority should be critical infrastructure: systems whose failure would cascade into catastrophic regional or broader national disruption. The government can provide legal protections, policy guidance and convening authority to solve this challenge. **The clearinghouse, established by Executive Order 14409 and launched as the "Gold Eagle Initiative" on July 14 2026, is designed to stage the release of open source patches to qualified parties before public disclosure in order to mitigate the risk of exploitation**. Public-private partnership is imperative since 50% to 85% of US critical infrastructure is privately owned and operated. Which organizations qualify as part of the Gold Eagle Initiative as critical infrastructure for purposes of addressing and prioritizing disclosure of patches remains to be clarified; such a designation could inform eligibility for the vulnerability clearinghouse, priority access to defensive tools and subsidized or shared access to defensive AI capabilities. Gold Eagle should be permanent and not temporary, and should be applied to lifeline sectors (communications, energy, water, transportation, banks and cloud service providers) given the new threat environment.

Domestic cyber-defense capability. Some cyber incidents could cascade across lifeline sectors and may be too large for any single firm to manage, particularly when nation-state actors are already targeting or maintaining access inside critical networks. One longer term policy option is a **dedicated, government-hosted domestic cyber-defense "strike force" capability with clearer authority to support lifeline sectors before and during serious incidents**. Such a capability would need to be distinct from a foreign intelligence mission, include appropriate civil liberties protections and operate with clear rules of engagement and strong oversight. It could perform more than an advisory role for private enterprise, providing direct operational support in high consequence scenarios.

Pre-release oversight of covered frontier models. The release of models with advanced cyber capabilities raises national security and critical infrastructure concerns. Executive Order 14409 issued in June 2026 creates a voluntary framework for federal national security risk assessments of "covered frontier models" before release. The government should clearly articulate a consistent framework for how the most advanced AI models should be screened for potential risks and rolled out to critical defenders, and eventually to broader release, rather than the current approach which has been more ad hoc.

Scaling AI-led defense adoption and modernization of weak links. Small and under-resourced operators often run end-of-life technology they cannot afford to replace. The problem is most acute in operational technology, where equipment stays in service for long periods and may be difficult or impossible to patch. Policy options include federal financing, targeted "rip-and-replace" incentives where obsolete systems create national-level risk (a "cash for clunkers" program for infrastructure instead of internal combustion engine vehicles), expanded CISA technical assistance and state and local cyber corps, and trusted-access programs to put defensive models within reach of designated, under-resourced operators.

International coordination. Cyber risk does not respect borders. The software supply chain is global, frontier AI development is increasingly international and adversaries route through infrastructure across jurisdictions. US coordination with allies on vulnerability disclosure, critical infrastructure resilience, AI model security, export controls, crisis communications and supply-chain resilience will be important to monitor. Fragmentation is a related risk: misaligned approaches could slow disclosure, complicate incident response, splinter maintainer communities, multiply poorly maintained code and expand the global attack surface.

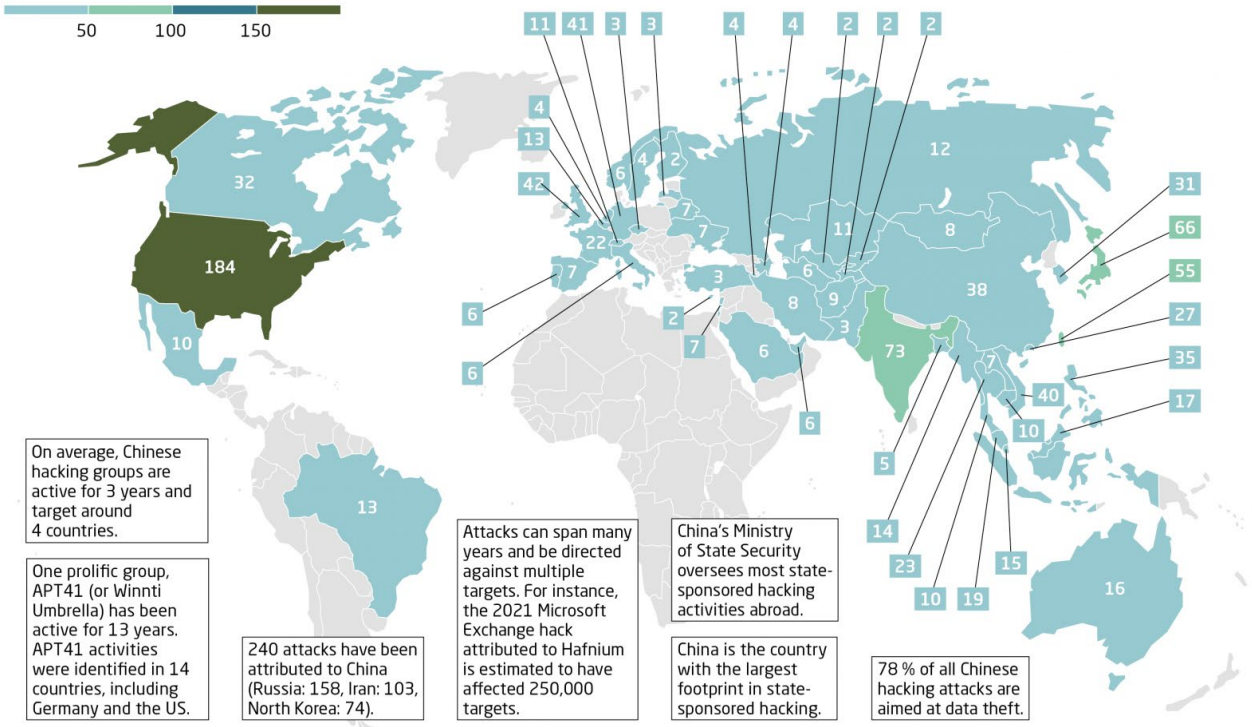
Bottom line

AI-enabled vulnerability discovery will increase the number of serious findings, shorten response timelines and expand the pool of actors capable of exploitation. The federal policy areas to watch are prioritization of the most important systems, faster collective defense, support for shared software and under-resourced operators, pre-release assessment of the most capable models and resilience planning for cases where patches do not arrive in time.

Appendix A: on cyberattacks, the US and China

From 2000-2023, China was responsible for 240 state-sponsored or state-affiliated cyberattacks, followed by Russia at 158 and Iran at 102. The US was on the receiving end of the highest number of such attacks.

Number of cyberattacks experienced by each country



Source: Mercator Institute for China Studies

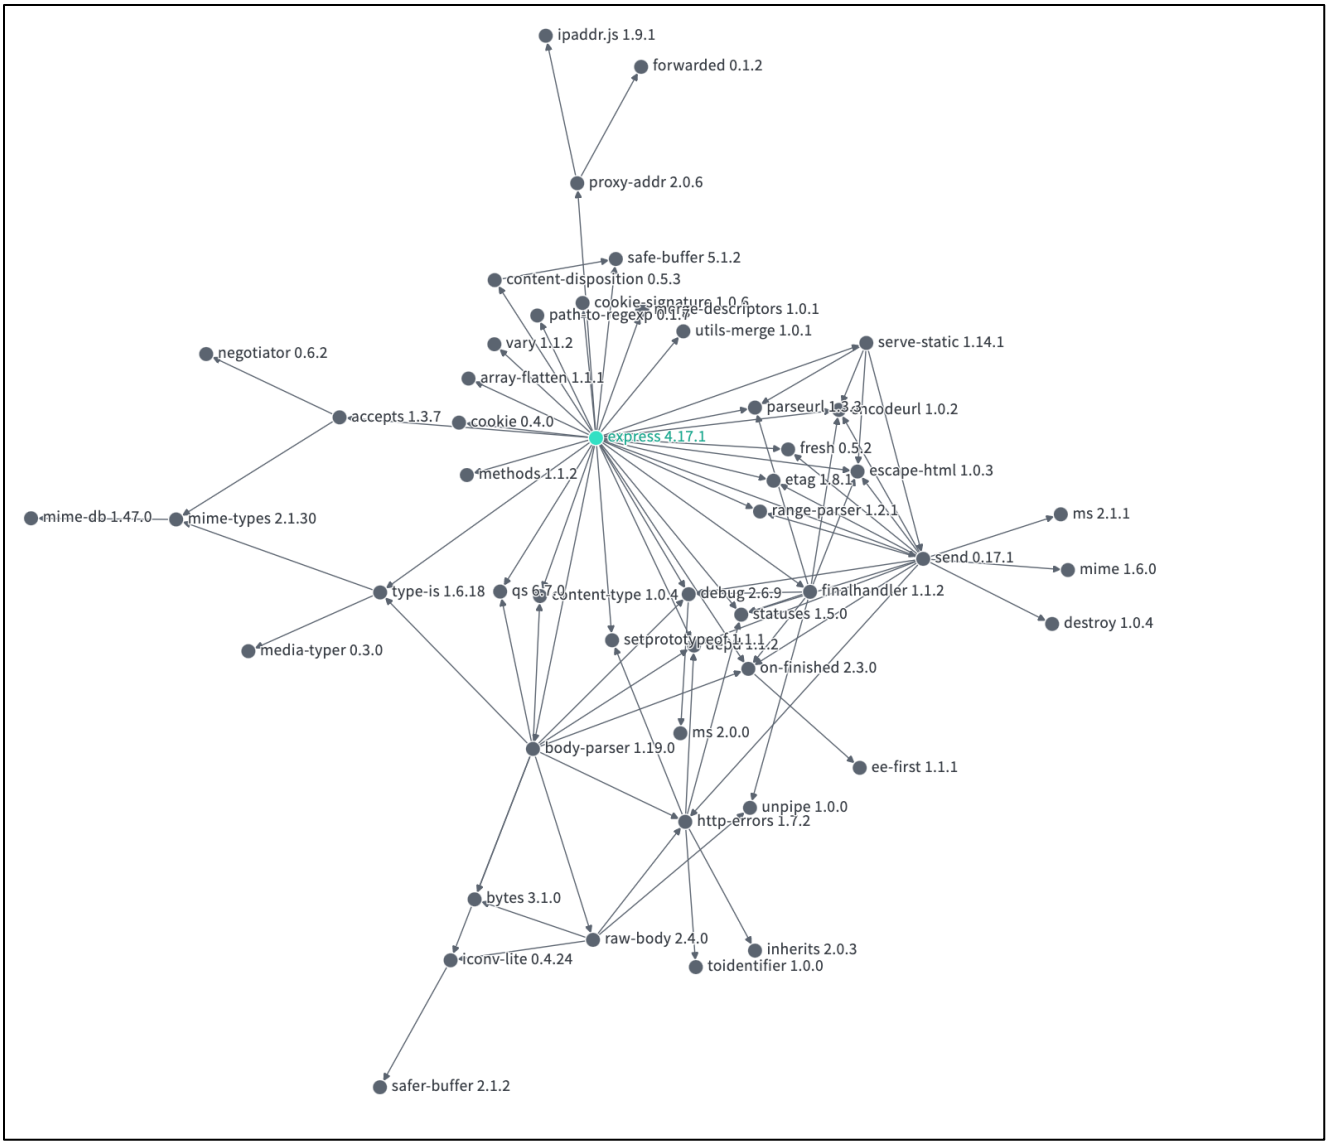
The table below outlines some notable vulnerability exploitations by China-nexus actors in 2024-2025.

APT name / aliases	Suspected affiliation	Primary mission profile	Key target sectors	Hallmark TTPs / malware
Volt Typhoon / Insidious Taurus, Bronze Silhouette, Voltzite	PRC State-Sponsored	Strategic pre-positioning in CNI for future disruptive/destructive attacks	Communications, energy, water, transportation (esp. in US & Guam)	Exclusive use of Living-off-the-Land (LOTL); KV Botnet (SOHO routers) for C2
APT41 / Double Dragon, Barium, Wicked Panda	PRC State-Sponsored (with parallel financial motive)	Dual-mission: State-sponsored espionage and financially motivated cybercrime	Espionage: healthcare, telecom, tech; crime: video gaming	Highly versatile; Exploits edge devices (Fortinet); Large malware arsenal (KeyPlug, Cobalt Strike)
APT40 / Leviathan, Bronze Mohawk, GINGHAM TYPHOON	Ministry of State Security (MSS)	Espionage supporting the Belt and Road Initiative and naval modernization	Maritime, defense, engineering, government (esp. in Southeast Asia)	Exploitation of SOHO routers for C2; Spear-phishing; Large toolset (SCANBOX)
Salt Typhoon	Ministry of State Security (MSS)	Espionage and intelligence collection focused on telecommunications providers	Telecommunications (Global, with focus on US)	Exploits Ivanti vulnerabilities; LoTL for lateral movement; Backdoors (GhostSpider, Masol RAT)

Source: "*An Analysis of China's Escalating Cyber Campaign Against Global Critical Infrastructure (2024-2025)*", FalconFeeds.io, August 21, 2025

Appendix B: an example of open source dependencies: Express 4.17

In 2021, Google launched the **Open Source Insights Project** which illustrates the complex web of open source interdependencies. The example below is for Express 4.17, open source code for building web applications and APIs. According to the NPM registry, there are over 16,000 software packages that rely on Express code.



IMPORTANT INFORMATION

This material is for information purposes only. The views, opinions, estimates and strategies expressed herein constitutes Michael Cembalest's judgment based on current market conditions and are subject to change without notice, and may differ from those expressed by other areas of JPMorgan Chase & Co. ("JPM"). **This information in no way constitutes J.P. Morgan Research and should not be treated as such.** Any companies referenced are shown for illustrative purposes only, and are not intended as a recommendation or endorsement by J.P. Morgan in this context.

GENERAL RISKS & CONSIDERATIONS Any views, strategies or products discussed in this material may not be appropriate for all individuals and are subject to risks. Investors may get back less than they invested, and **past performance is not a reliable indicator of future result**s. Asset allocation/diversification does not guarantee a profit or protect against loss. Nothing in this material should be relied upon in isolation for the purpose of making an investment decision.

NON-RELIANCE Certain information contained in this material is believed to be reliable; however, JPM does not represent or warrant its accuracy, reliability or completeness, or accept any liability for any loss or damage (whether direct or indirect) arising out of the use of all or any part of this material. No representation or warranty should be made with regard to any computations, graphs, tables, diagrams or commentary in this material, which are provided for illustration/ reference purposes only. Any projected results and risks are based solely on hypothetical examples cited, and actual results and risks will vary depending on specific circumstances. Forward-looking statements should not be considered as guarantees or predictions of future events. Nothing in this document shall be construed as giving rise to any duty of care owed to, or advisory relationship with, you or any third party. Nothing in this document shall be regarded as an offer, solicitation, recommendation or advice (whether financial, accounting, legal, tax or other) given by J.P. Morgan and/or its officers or employees. J.P. Morgan and its affiliates and employees do not provide tax, legal or accounting advice. You should consult your own tax, legal and accounting advisors before engaging in any financial transactions.

For J.P. Morgan Asset Management Clients:

J.P. Morgan Asset Management is the brand for the asset management business of JPMorgan Chase & Co. and its affiliates worldwide.

To the extent permitted by applicable law, we may record telephone calls and monitor electronic communications to comply with our legal and regulatory obligations and internal policies. Personal data will be collected, stored and processed by J.P. Morgan Asset Management in accordance with our privacy policies at https://am.jpmorgan.com/global/privacy.

ACCESSIBILITY
For U.S. only: If you are a person with a disability and need additional support in viewing the material, please call us at 1-800-343-1113 for assistance.

This communication is issued by the following entities: In the United States, by J.P. Morgan Investment Management Inc. or J.P. Morgan Alternative Asset Management, Inc., both regulated by the Securities and Exchange Commission; in Latin America, for intended recipients' use only, by local J.P. Morgan entities, as the case may be.; in Canada, for institutional clients' use only, by JPMorgan Asset Management (Canada) Inc., which is a registered Portfolio Manager and Exempt Market Dealer in all Canadian provinces and territories except the Yukon and is also registered as an Investment Fund Manager in British Columbia, Ontario, Quebec and Newfoundland and Labrador. In the United Kingdom, by JPMorgan Asset Management (UK) Limited, which is authorized and regulated by the Financial Conduct Authority; in other European jurisdictions, by JPMorgan Asset Management (Europe) S.à r.l. In Asia Pacific ("APAC"), by the following issuing entities and in the respective jurisdictions in which they are primarily regulated: JPMorgan Asset Management (Asia Pacific) Limited, or JPMorgan Funds (Asia) Limited, or JPMorgan Asset Management Real Assets (Asia) Limited, each of which is regulated by the Securities and Futures Commission of Hong Kong; JPMorgan Asset Management (Singapore) Limited (Co. Reg. No. 197601586K), which this advertisement or publication has not been reviewed by the Monetary Authority of Singapore; JPMorgan Asset Management (Taiwan) Limited; JPMorgan Asset Management (Japan) Limited, which is a member of the Investment Trusts Association, Japan, the Japan Investment Advisers Association, Type II Financial Instruments Firms Association and the Japan Securities Dealers Association and is regulated by the Financial Services Agency (registration number "Kanto Local Finance Bureau (Financial Instruments Firm) No. 330"); in Australia, to wholesale clients only as defined in section 761A and 761G of the Corporations Act 2001 (Commonwealth), by JPMorgan Asset Management (Australia) Limited (ABN 55143832080) (AFSL 376919). For all other markets in APAC, to intended recipients only.

For J.P. Morgan Private Bank Clients:

ACCESSIBILITY
J.P. Morgan is committed to making our products and services accessible to meet the financial services needs of all our clients. Please direct any accessibility issues to the Private Bank Client Service Center at 1-866-265-1727

LEGAL ENTITY, BRAND & REGULATORY INFORMATION
In the **United States**, **JPMorgan Chase Bank, N.A.** and its affiliates (collectively "**JPMCB**") offer investment products, which may include bank managed investment accounts and custody, as part of its trust and fiduciary services. Other investment products and services, such as brokerage and advisory accounts, are offered through **J.P. Morgan Securities LLC ("JPMS")**, a member of FINRA and SIPC. JPMCB and JPMS are affiliated companies under the common control of JPM.
In **Germany**, this material is issued by **J.P. Morgan SE**, with its registered office at Taunustor 1 (TaunusTurm), 60310 Frankfurt am Main, Germany, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB). In **Luxembourg**, this material is issued by **J.P. Morgan SE – Luxembourg Branch**, with registered office at European Bank and Business Centre, 6 route de Treves, L-2633, Senningerberg, Luxembourg, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Luxembourg Branch is also supervised by the Commission de Surveillance du Secteur Financier (CSSF); registered under R.C.S Luxembourg B255938. In the **United Kingdom**, this material is issued by **J.P. Morgan SE – London Branch**, registered office at 25 Bank Street, Canary Wharf, London E14 5JP, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – London Branch is also supervised by the Financial Conduct Authority and Prudential Regulation Authority. In **Spain**, this material is distributed by **J.P. Morgan SE, Sucursal en España**, with registered office at Paseo de la Castellana, 31, 28046 Madrid, Spain, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE, Sucursal en España is also supervised by the Spanish Securities Market Commission (CNMV); registered with Bank of Spain as a branch of J.P. Morgan SE under code 1567. In **Italy**, this material is distributed by **J.P. Morgan SE – Milan Branch**, with its registered office at Via Cordusio, n.3, Milan 20123, Italy, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Milan Branch is also supervised by Bank of Italy and the Commissione Nazionale per le Società e la Borsa (CONSOB); registered with Bank of Italy as a branch of J.P. Morgan SE under code 8076; Milan Chamber of Commerce Registered Number: REA MI 2536325. In the **Netherlands**, this material is distributed by **J.P. Morgan SE – Amsterdam Branch**, with registered office at World Trade Centre, Tower B, Strawinskylaan

1135, 1077 XX, Amsterdam, The Netherlands, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Amsterdam Branch is also supervised by De Nederlandsche Bank (DNB) and the Autoriteit Financiële Markten (AFM) in the Netherlands. Registered with the Kamer van Koophandel as a branch of J.P. Morgan SE under registration number 72610220. In **Denmark**, this material is distributed by **J.P. Morgan SE – Copenhagen Branch, filial af J.P. Morgan SE, Tyskland**, with registered office at Kalvebod Brygge 39-41, 1560 København V, Denmark, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Copenhagen Branch, filial af J.P. Morgan SE, Tyskland is also supervised by Finanstilsynet (Danish FSA) and is registered with Finanstilsynet as a branch of J.P. Morgan SE under code 29010. In **Sweden**, this material is distributed by **J.P. Morgan SE – Stockholm Bankfilial**, with registered office at Hamngatan 15, Stockholm, 11147, Sweden, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Stockholm Bankfilial is also supervised by Finansinspektionen (Swedish FSA); registered with Finansinspektionen as a branch of J.P. Morgan SE. In **Belgium**, this material is distributed by **J.P. Morgan SE – Brussels Branch** with registered office at 35 Boulevard du Régent, 1000, Brussels, Belgium, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE Brussels Branch is also supervised by the National Bank of Belgium (NBB) and the Financial Services and Markets Authority (FSMA) in Belgium; registered with the NBB under registration number 0715.622.844. In **Greece**, this material is distributed by **J.P. Morgan SE – Athens Branch**, with its registered office at 3 Haritos Street, Athens, 10675, Greece, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Athens Branch is also supervised by Bank of Greece; registered with Bank of Greece as a branch of J.P. Morgan SE under code 124; Athens Chamber of Commerce Registered Number 158683760001; VAT Number 99676577. In **France**, this material is distributed by **J.P. Morgan SE – Paris Branch**, with its registered office at 14, Place Vendôme 75001 Paris, France, authorized by the Bundesanstaltfür Finanzdienstleistungsaufsicht(BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB) under code 842 422 972; J.P. Morgan SE – Paris Branch is also supervised by the French banking authorities the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and the Autorité des Marchés Financiers (AMF). In **Switzerland**, this material is distributed by **J.P. Morgan (Suisse) SA**, with registered address at rue du Rhône, 35, 1204, Geneva, Switzerland, which is authorised and supervised by the Swiss Financial Market Supervisory Authority (FINMA) as a bank and a securities dealer in Switzerland.

In **Hong Kong**, this material is distributed by **JPMCB, Hong Kong branch**. JPMCB, Hong Kong branch is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission of Hong Kong. In Hong Kong, we will cease to use your personal data for our marketing purposes without charge if you so request. In **Singapore**, this material is distributed by **JPMCB, Singapore branch**. JPMCB, Singapore branch is regulated by the Monetary Authority of Singapore. Dealing and advisory services and discretionary investment management services are provided to you by JPMCB, Hong Kong/Singapore branch (as notified to you). Banking and custody services are provided to you by JPMCB Singapore Branch. The contents of this document have not been reviewed by any regulatory authority in Hong Kong, Singapore or any other jurisdictions. You are advised to exercise caution in relation to this document. If you are in any doubt about any of the contents of this document, you should obtain independent professional advice. For materials which constitute product advertisement under the Securities and Futures Act and the Financial Advisers Act, this advertisement has not been reviewed by the Monetary Authority of Singapore. JPMorgan Chase Bank, N.A., a national banking association chartered under the laws of the United States, and as a body corporate, its shareholder's liability is limited.

With respect to countries in **Latin America**, the distribution of this material may be restricted in certain jurisdictions.

*Issued in **Australia** by JPMorgan Chase Bank, N.A. (ABN 43 074 112 011/AFS Licence No: 238367) and J.P. Morgan Securities LLC (ARBN 109293610).*

References to "J.P. Morgan" are to JPM, its subsidiaries and affiliates worldwide. "J.P. Morgan Private Bank" is the brand name for the private banking business conducted by JPM. This material is intended for your personal use and should not be circulated to or used by any other person, or duplicated for non-personal use, without our permission. If you have any questions or no longer wish to receive these communications, please contact your J.P. Morgan team.

© 2026 JPMorgan Chase & Co. All rights reserved.

J.P.Morgan