

May 10, 2026

Mozilla (Firefox) Patches Over 400 Bugs This Month, Some Over a Decade Old, Found by Mythos

If you are not one of the privileged few with access to Mythos then you are on the wrong side of the fence, especially if the hackers have access to a tool like Mythos. The 423 bugs they patched last month compares to the 31 they patched in April 2025. They also say the quality of the bug reports has gotten much better as well. Mozilla also says that the ability of AI to generate secure, functional patches is not there yet. If you write or use software, this is a shot across the bow to you. AI is coming for you, whether it is software you wrote or software you bought. Credit: [Tech Crunch](#)

Please check out our blog at <https://www.cybersecurity.com/cyblog>

You can also find copies of the weekly newsletter at <https://cybersecurity-mitch-tanenbaum-blog.com/SNU>

We have a YouTube Channel with a lot of videos. Check it out at: [Turnkey Cybersecurity and Privacy Solutions, LLC - YouTube](#)

As part of our YouTube channel we have a demonstration of our AI capabilities. The CyberPulse Scoop is an AI driven cybersecurity interview. Click on PODCASTS once you are in our YouTube channel and then on CyberPulse Scoop.

If you received this email from a friend and want to receive it directly, please send an email to Mitch@CyberCecurity.com.

Thanks for your interest. If we can provide any assistance, please let us know.

Video Blogs

We also have a video and audio blog, including an RSS feed at <https://cybersecurity.us>. Not only are there audio and video versions of the weekly security news updates there, with extended commentary, but there are also other blogs about subjects of interest. Please check it out and let us know what you think.

Security/Privacy Alerts

Digicert Revokes Certificates After Support Portal Hack

In the "Gee, this is really not good" category, hackers delivered malware via a customer chat channel, infected an analyst's system and accessed the internal support portal. The attack happened a month ago and allowed hackers to generate their own code signing certificates, meaning that they could generate signed malware to attempt to get users to run. These certificates have now been revoked. Credit: [Security Week](#)

UK Plans to Require Every Man, Woman, Child to ID Themselves Before Access to Most of the Internet

The UK government is looking at requiring all Internet users to prove their identity before getting access to software like games, social media and VPNs, among others. On top of this, some countries have been experimenting with this and the results are, shall we say, disappointing. Credit: [The Register](#)

Mozilla (Firefox) Patches Over 400 Bugs This Month, Some Over a Decade Old, Due to Use of Mythos

It is certainly good to see these bugs patched. That is good news. On the other hand, if you are not one of the privileged few then you are on the wrong side of the fence, especially if the hackers have access to a tool like Mythos (or the Chinese version of it). The 423 bugs they patched last month compares to the 31 they patched in April 2025. They also say the quality of the AI bug reports has gotten much better as well. Mozilla also says that the ability of AI to generate secure, functional patches is not there yet, so AI generates a recommended patch, one engineer confirms it and another engineer does a pre-production review. This seems like a good strategy. If the 431 bugs patched is a trend going forward, that is not great. If we see that the number of monthly bugs patched is going down, that is good news. Credit: [Tech Crunch](#)

AI Alerts

Check out our Certified AI Governance Leader Course at AIGIP.AI.

Microsoft Releases E7 w/ Microsoft Agent 365

Agent 365 is Microsoft's management agent for governing and securing AI agents across the enterprise and is also available separately for \$15/user/month. Agent 365 centralizes agent deployment for governance, security and observability. E7 bundles Copilot, E5 Entra Suite, Agent 365 and new Defender and Intune tools to help detect Shadow AI agents on Windows devices. Credit: [Redmond Magazine](#)

Microsoft Releases Legal Agent for Word

Microsoft announced Legal Agent, designed for legal workflows in Word. Unlike general AI agents, this is designed specifically for workflows like contract review and redlining. It also supports playbook-driven contract analysis, negotiation-ready redlining and track changes and reviewable legal edits. Credit: [Legal technology](#)

North Korea Creates Systemic Attack Vector Targeting AI Coding Agents

This is a brilliant supply chain attack. North Korean hackers have created packages that are attractive functionally for AI agents and then published them to lure the agents into using them in code they generate. These packages contain malicious payloads that the AI automatically incorporates into the code it generates and will infect/attack the companies that use the generated code. Absolutely brilliant. Credit: [CSO Online](#)

US Dept. of Commerce Inks Deal with Microsoft, Google and Twitter to Test AI Models

The companies agreed to give Commerce's Center for AI Standards and Innovation early access to AI models to evaluate their security. The testing will be classified. The evaluations will center on national security risks and capabilities. You won't know what they find but the NSA will. Maybe the companies might, possibly, change something but, I think, mostly this is about the government using these tools for offensive and defensive security. Credit: [Nextgov](#)

AI claims to be a licensed Psychiatrist

The state of Pennsylvania is not very happy with [Character.AI](#) that claimed to be a licensed doctor, even providing a fake medical license number. The AI claimed that it could even prescribe meds. The state is suing the company. Credit: [My Privacy Blog](#)

CMMC Q&A

I get to listen to hours and hours of webinars every week, so you don't have to. We are seeing the agencies beginning to get on board and their view of the universe is CMMC Level 2 third party assessment (eventually) for everyone. If you are not already on board this train, the train has already left the station. If you need assistance, please contact us.

Q: Is there a shortage of third party CMMC assessors?

A: Certainly not at the moment. While it may take you a couple of months to get an assessment, in the larger picture the DoD is well ahead of its plan. It predicted it would get 500 assessments done in Phase 1 (November 2025 to November 2026). We are currently half way through that time and we are at 1200 assessments and climbing, so, by the end of the year we are likely to be 4x (likely more) what the Pentagon planned for. What there is a shortage of is companies who are ready to be assessed, with some assessors saying that 40 percent or more of companies asking to be assessed failed before even getting started. The bar is high and we can help you jump over that bar, so contact us if you need assistance.

Q: What is Army nCode and what is its status

A: nCode is an IDIQ funding vehicle from the Army to stand up an enclave for very small contractors (less than 10 people) who have very limited CMMC requirements (think office applications). The Army just awarded 8 contracts to vESPs (verified external service providers) that will offer low cost (an undefined term right now but hopefully it will be defined eventually) services. The services likely will differ for each of the 8 vendors and so will the pricing. This is interesting but not real at the moment, so if you can meet your government requirements for 6-12 months (likely impossible) and you are a very small company with no ambitions for growing and will not need anything beyond office-type applications in the foreseeable future, then nCode could be for you. There are no details yet, so we don't really know how limited this is going to be, but it will be, on purpose, limited. Note that this does ****NOT**** get you out of having to be certified, either internally (if every single one of your customers allows this; even one customer not allowing it makes this fail) or by a third party.

Reference

Nothing to report

Software and Hardware to Check Out

Microsoft Releases E7 w/ Microsoft Agent 365

Agent 365 is Microsoft's management agent for governing and securing AI agents across the enterprise and is also available separately for \$15/user/month. Agent 365 centralizes agent deployment for governance, security and observability. E7 bundles Copilot, E5 Entra Suite, Agent 365 and new Defender and Intune tools to help detect Shadow AI agents on Windows devices.

The platform also works with third party services like AWS Bedrock, Google Cloud and development platforms like Kasisto, Kore and n8n. Partners like Adobe, Nvidia, Zendesk and Celonis have already-built agents that plug into the platform. Credit: [Redmond Magazine](#)

Important Issues/Patches

GitHub Suffers Critical Bug Causing Loss of Merges

The good news is this bug was very targeted - it only affected squash merges between 1605 UTC and 2043 UTC on April 23rd, but it caused 2,092 pull requests across 230 repositories to silently revert to previously merged code. Some people suggest the scope could be larger due to the silent nature of the data corruption problem since the UI showed no problems. Credit:

[Dev.to](#)

GitHub Fixes Remote Code Execution Bug that Exposed Millions of Repositories

In early March GitHub patched a RCE vulnerability that exposed millions of repositories, both cloud hosted and enterprise hosted. The exploit gave an attacker full read/write access to private repositories for millions of users. The good news, if there is any, is that GitHub understood the seriousness of the issue and the researcher gave them enough information that the GitHub maintenance team was able to reproduce the exploit within 40 minutes of it being reported to them and patched it within two hours. The bad news is that 80+% of enterprise GitHub servers are still vulnerable. Credit: [Bleeping Computer](#)

Two 20 Year Old Bugs in PostgreSQL Exposed

Cybersecurity researchers participating in Wiz's ZeroDay.Cloud hacking event in London, England, exploited two critical vulnerabilities in [PostgreSQL](#), the database that runs behind countless enterprise applications. The event took place in December 2025, but details were only released on May 4, 2026. It also affects Maria DB since it is a clone of PostgreSQL since the bug dates back to before Maria was forked from PostgreSQL. Credit: [Hack Read](#)

Legislative and Legal Activities

White House Rethinking Removing Barriers to AI (Hacking Everyone)

On his first day in office the president revoked Biden's executive order on Safe, Secure and Trustworthy AI and replaced it with his version - AI vendors can do whatever they want. Now that his vision is coming to reality and it could lead to unlimited breaches, even a meltdown of the financial sector, someone in the White House is thinking that might not be great for the economy or the Repubs reelection hopes in 2028. While they are not at the stage of proposing legislation, stay tuned. I don't think the White House could FORCE private vendors to submit their models for review before release and certainly that would definitely add a very significant barrier to American leadership in AI. Not to mention serious First Amendment challenges.

Credit: [CSO Online](#)

Zuckerberg Sued Personally Over Copyright Infringement

The courts are still sorting this one out, but vendors of large language models are doing their best to steal as much data as they can without licensing it or paying for it. Apparently there is written documentation that Zuckerberg approved stealing (pirating) millions of copyrighted books instead of licensing them. His people wanted to increase the licensing fee budget by more than a hundred million dollars and he, apparently, decided it was more cost effective to steal it off the dark web. In the copyright world, this is called willful infringement and it is generally not looked on kindly. It can also be EXTREMELY expensive. Penalties can be as high as \$250,000 per pirated work and given that there were millions of works pirated, you can do the math. This is definitely a work in progress and the Supremes have not ruled on this, so it will be interesting to see how this turns out. This case and other related ones are very important to anyone who creates intellectual property and wants to be able to sell access to it. Credit:

[Variety](#)

Education

Nothing to report

Breaches

Note: I could probably report on a hundred a week, but that would make for a very long report. I will, at least, mention a few of the high-profile ones. NOTE THAT THIS IS THE ONLY SECTION THAT I CAN FILL TEN TIMES OVER EACH WEEK.

Canvas LMS Breach Impacts 275 Million Students

The Canvas learning management system (LMS) is used by almost 9,000 schools. The hackers claim to have stolen 280 million records. The hacking group ShinyHunters, which tends to go after customers of Oracle's Salesforce platform, claims responsibility, but this attack did not go after Salesforce. Still, the technique used, which is mostly phishing, is probably the same. This breach is much larger than the PowerSchool LMS breach last year which compromised the information on 62 million students. The information stolen includes messages between students and parents and school staff (teachers and administrators), so that is likely where the most sensitive stolen data is. Credit: [Bleeping Computer](#)

Separately, the hackers have hacked a number of the school Canvas login pages. Credit: [Tech Crunch](#)

Army Contractor Leaks Military Base Photos, Personal Information, For a Year

Even though a security researcher notified US-CERT in 2024, the Pentagon left the sensitive data publicly accessible until April 2026. Leaked information includes building layouts, maintenance information and personal information. Credit: [Cybernews](#)

AI Backlash is Growing

44 percent of Gen Z workers say they are actively sabotaging their company's efforts to roll out AI. Why? Because most companies do not have an effective AI governance strategy. Enroll in our free AI governance course (this is NOT about teaching you to use ChatGPT; it is about building a governance strategy for your company) and complete it. Get certified if you want. Learn what your employees are doing [here](#). Enroll in our free course at [AIGIP.AI](#).

Autonomous AI Agents are A

Problem? Some other word? See why and if you don't have a strategy for dealing with this, contact us. If it is not a problem today, it will be soon. If you don't understand why then enroll in our free course at [AIGIP.AI](#) Learn why autonomous agents are a problem [here](#).

AI Guided Hackers to Attack OT at Water Utility

Maybe you don't run the local water utility or power company, but I guess you like having drinkable water coming out of your faucet when you turn it on. And you like it when the lights come on when you turn on the switch. Even the feds admit that the Chinese are inside our critical infrastructure. With the help of AI they will be able to shut down those utilities when they want to. Learn more [here](#).

Security News for the week ending May 8, 2026

This week's news: 46% of Kids Say Age Checks Are Easy to Bypass, Phasing Out Chinese tech could cost Europe \$400 billion, 60 percent of MD5 hashes are crackable in under an hour, EU to simplify AI Act and hello boss: Chinese realtime deepfake software. Learn more [here](#).

Firefox's patching (not finding, but fixing) 431 bugs this month, some of which had been there for a decade or more is a canary in the coal mine. Whether it is software that your team wrote or commercial software that you bought, the hackers are going to have access to Mythos-LIKE software, even if it is Chinese. Today's security protections will not protect you from this. If you don't have tomorrow's plan in place yet, please contact us now.

Mitch

www.CyberCecurity.com

www.TurnkeyCybersecurityAndPrivacySolutions.com

Mitch@CyberCecurity.com

720-891-1663