

# Security Now! #1079 - 05-19-26

## Daybreak and Codename MDASH

### This week on Security Now!

- Microsoft rethinks Edge's "intended behavior" after it gets press.
- Chaotic Eclipse hacker strikes again with a Bitlocker bypass.
- Google's threat analysis group documents malicious AI use.
- Canada hasn't learned the lessons of the EU and the UK.
- AI chatbots may be far more addictive than social media.
- Project: Hail Mary now available to stream.
- An apparently-serious zero-point quantum vacuum energy source.
- A bit of listener feedback.
- OpenAI's & Microsoft's vulnerability discovery systems.

**Worries over AI surpassing us may be overblown because AI has been trained on human output.**



## Security News

### It is "Intended Behavior" only until it gets media attention

Last week, we noted the discovery, reporting and widespread confirmation that Microsoft's Edge browser was storing all of its users' passwords in RAM where they were easily discoverable and exfiltrateable en masse. This data included the URLs and the usernames and passwords required to login to every website whose data was present and, presumably, where no other authentication factor would be required.

I'll pause to note that this is a perfect example of the reason why, if one is going to go to the trouble of having additional factors for authentication security, it's nuts to store that additional authentication information with the same single provider. Listeners have asked whether it is safe to store their one-time password secrets in the same password manager as their usernames and passwords. They want me to say "yes", since it's so convenient to extend a password manager's capabilities to include responding to the query for 6-digit OTP tokens. I really do get it and I understand the temptation. So I'll just say that I have never done that, and I never would nor will do that. The entire point here is separation and redundancy which is completely lost when all of the eggs are stored in a single basket. I use "OTP Auth" on my separate iPhone. The good news is that most sites have become smarter about avoiding the needless prompting for one-time password tokens. Whereas a financial institution or the government might reasonably insist upon the provision of a one-time password with every login, many other less sensitive sites that have been configured to require a one-time password will relax their need when the browser being used already carries a previously-valid login cookie which indicates that the browser was previously logged into that site. This is the newer "we recognize you on this computer" messaging that we're seeing more often. And that's good, since we want bad guys, who will not have that browser cookie, to be forced to come up with that additional authentication factor, whereas we don't want it to be overly burdensome for regular users who want that added safely without the overboard hassle.

So it should be clear that if a single source of breached authentication data also contained the one-time password secrets then all of that extra protection the user hoped to obtain would be lost. I know it's more hassle, but multiple sources of authentication should always be as heterogeneous as possible and practical.

Okay. Stepping down from my soapbox and back to Microsoft and Edge. Last week we noted that Microsoft's disappointing by predictable response to questioning about their in-the-clear storage of the users' authentication data was that it was "intended behavior." The SANS Security Institute write: *"Microsoft classifies this as "intended behaviour." I'm not sure what manager or lawyer decided that, hopefully it wasn't anyone in their security team."* Since I titled this first bit of news "Intended behavior" only until it gets media attention you can guess what comes next.

BleepingComputer provides the details and background, writing last Friday:

*Microsoft is updating the Edge web browser to ensure it no longer loads saved passwords into process memory in clear text at startup after previously stating it was "by design." This behavior was disclosed on May 4 by a security researcher Tom Rønning, who demonstrated that all credentials stored in the Edge built-in password manager were decrypted on launch*

*and kept in memory even when not in use. Rønning also released a proof-of-concept (PoC) tool that would allow attackers with Administrator privileges to dump passwords from **other** users' Edge processes (without admin privileges, the PoC only allows accessing Edge processes launched by the same user). He said he reported the issue to Microsoft and was told the behavior was "by design" before he publicly disclosed it.*

I'll note that this is an interesting wrinkle on the "responsible disclosure" principle. You tell someone responsible, like Microsoft, in confidence, about some clearly bad behavior you've just discovered in one of their highly security critical flagship products and you're quite clearly told "Yeah, that's right. That's what we want, so that's the way it is." At that point, no one is going to fault you for letting the rest of the world know what you've found and that you were told to buzz off. BleepingComputer quotes the discoverer:

*"Edge is the only Chromium-based browser I've tested that behaves this way. By contrast, Chrome uses a design that makes it far harder for attackers to extract saved passwords by simply reading process memory."*

*While it initially refused to address the issue, telling BleepingComputer at the time that "this is an expected feature of the application," Microsoft announced on Wednesday that future versions of Edge will no longer load saved passwords into memory on startup, even though the reported scenario falls within the expected existing threat model (which excludes attacks where an adversary already has administrative control of a device).*

*Microsoft Edge Security Lead Gareth Evans said: "This defense-in-depth change will come to every supported version of Edge (Stable, Beta, Dev, Canary, and the Extended Stable channel our enterprise customers run), and we're prioritizing the rollout. With our commitment to the Secure Future Initiative and customer feedback, we are taking a broader view. That means looking not only at whether something meets the bar for a security issue, but also at where we can reduce exposure through defense-in-depth improvements. In this case, reducing the exposure of passwords in memory is a practical step in that direction."*

*[BleepingComputer reports that] The fix is already live in the Edge Canary channel and will be included in the next update for all supported Edge releases (build 148 and newer). Last year, Microsoft also introduced a new Edge security feature to protect users against malicious extensions sideloaded into the web browser, and restricted access to Edge's Internet Explorer mode after hackers began leveraging zero-day exploits in the Chakra JavaScript engine to access target devices.*

Okay. So first, while writing this last Saturday I immediately fired up Edge to check its Help/About and watched it quickly updating itself to build 148. So that fix was, indeed, quickly pushed out.

But the point Microsoft made about the threat model governing Edge's design was important, reasonable and worthy of additional attention. BleepingComputer wrote: *"Microsoft announced on Wednesday that future versions of Edge will no longer load saved passwords into memory on startup, even though the reported scenario falls within the expected existing threat model (which excludes attacks where an adversary already has administrative control of a device)"* In other words, they're saying "we are going to change this behavior even though the scenario Tom Ronning discovered, where all username and password authentication was being needlessly

pre-loaded into RAM, falls within the expected existing threat model." Before I defend Microsoft's response I'll take exception to their use of the term "administrative control of a device." As was noted, administrative control is explicitly not required. Administrative control allows malware to obtain the usernames and passwords of ALL of a system's users who might be logged in at the time with Edge running. But malware running in a non-admin account can still access all of its own user's in-RAM Edge authentication.

But let's focus upon the intent behind Microsoft's defensive position. The concept and deliberate design of formal threat models is perhaps the most important advance in our understanding and practice of security. We saw a lot of that during last week's deep dive into DigiCert's internal security architecture. Just the fact that an "architecture" is something "security" can have, represents a tremendous advance in the state-of-the-art of our understanding of how to provide protection.

So in this case, Microsoft is essentially saying: "We recognize that once an attacker has taken up residence inside a system – by whatever means – our ability to limit the damage that could be done is severely limited by the tradeoffs we've had to make in the name of practical usability. A perfect example is User Account Control. I may refuse to store my one-time password secrets in my password manager, but the first thing I do when setting up a new Windows machine, before I totally lose my mind, is completely disable UAC. Having that thing constantly darkening my doorway – I mean, my screen – and popping up to get my permission when I want to do perfectly safe things – the consequences of which I fully understand – is not offering a value proposition that works for me. Have I sacrificed some security by disabling that nagging nanny? You betcha. Yep. No doubt. But my sanity is important to me.

Microsoft is in an impossible position. I'm entirely sympathetic. Windows is being used by people who will follow commands provided to them by some random page on the Internet, instructing them to blindly paste and run a command they could not possibly understand, even if they could see it. So how is Microsoft supposed to protect such users from themselves when an increasingly hostile world wants to attack them?

So on the one hand, Microsoft's position that there can be no true protection from bad guys who have already gotten into one's PC is accurate and defensible. In fact, in a minute or two we're going to examine what's been dubbed "The Bitlocker Bypass." It's a perfect case in point about the nature of local compromises and security boundaries.

But the other point Microsoft made, quoting the phrase "defense in depth" refers to another of the crucial advances that have been made in our contemporary understanding of security. When a castle was surrounded by a piranha-filled moat, attackers could bring a boat. But when the outside of that moat is surrounded by a tall fence it would be difficult to get the boat to the moat. By the way, "defense in depth" is also exactly storing all authentication factors in a single place should always be avoided when storing them in multiple separate places is possible.

The bottom line is that the attention drawn to Edge's needless exposure of its usernames and passwords revealed ... that the exposure was needless. As we saw, none of the other Chromium-based web browsers behave so cavalierly with their users' most important secrets. Every one of the others goes to the time and trouble to protect them. Now Edge does too. So that's good.

## The Bitlocker Bypass

While we're on the topic of Microsoft I wanted to make sure that everyone knew about the recent discovery – with a published proof of concept – of a local bypass attack on Microsoft's proprietary Bitlocker drive encryption.

The source and the apparently deliberate timing of the disclosure of this latest significant Windows vulnerability is interesting because it was publicly released last week on the 13th, the day after this month's Patch Tuesday. And who released it? None other than the hacker Chaotic Eclipse with his Nightmare-Eclipse Github account. This is the individual we talked about recently who is extremely perturbed by Microsoft's handling of him and his disclosures. Recall that he appears to accuse and blame Microsoft for deliberately and knowingly ruining his life. In retaliation for that perceived sleight he had previously disclosed the BlueHammer and RedSun local privilege escalation vulnerabilities as 0-day flaws, both which began being exploited in the wild shortly after being publicly disclosed.

So now, Chaotic Eclipse is back, publishing two new exploits with proofs for two new unpatched vulnerabilities named YellowKey and GreenPlasma. They are respectively, a BitLocker bypass and a privilege-escalation. He describes the BitLocker bypass issue as functioning like a backdoor because the vulnerable component is present only in the Windows Recovery Environment (WinRE), used as a utility host, often to repair boot-related problems with Windows. This "Chaotic Eclipse" individual remains miffed at Microsoft and so has published guidance on how to exploit them. Finally, this researcher has promised what he described as "a big surprise" for the next Patch Tuesday.

Security researcher Kevin Beaumont who posts as "GossiTheDog" has independently confirmed the functioning of the "YellowKey" Bitlocker bypass. Kevin first posted over on Mastodon: "*So I've just had a quick play with this and yes, it works. Essentially BitLocker has a backdoor. Mitigation = BitLocker PIN and BIOS password lock.*" Of course, a BIOS password lock is a pain in the butt, but for high-risk scenarios where local access with rebooting might be possible, it might be the strongest and quickest cure until Microsoft arranges to fix this.

Kevin followed his first Mastodon posting with a thread of posts which said:

*I think my prior toot on NightmareEclipse auto deleted so to make a perm one - it isn't me. I suspect it's somebody who used to work at MSFT, who departed after my era. For anybody looking at this, testing showed two things:*

*- TPM unlocked the storage, it provides a login bypass, as you're dumped as SYSTEM prior to Windows Hello or password login. BitLocker operates without a PIN by default, so it's basically a big gap, it's unclear how this code made it into the production version of Windows. I should point out I've only tested with one version of Windows 11 - maybe the scope is smaller. Will Dormann and I have both recreated the BitLocker backdoor -er- vulnerability.*

So what's the story here? BleepingComputer's headline was "*Windows BitLocker zero-day gives access to protected drives, PoC released*". Since we already have a lot about the background, I'm going to skip to their description of the trouble and excerpt some good bits. They write:

*The researcher says that YellowKey is a BitLocker bypass that affects Windows 11 and Windows Server 2022/2025. It involves placing specially crafted 'FsTx' files on a USB drive or EFI partition, rebooting into WinRE, and triggering a shell by holding down the CTRL key. The BitLocker bypass should also work without USB storage by copying the files to the EFI partition on the target drive. According to Chaotic Eclipse, the spawned shell gains unrestricted access to the storage volume protected by BitLocker.*

*Independent security researcher Kevin Beaumont confirmed that the YellowKey exploit is valid and agreed that BitLocker has a backdoor. He recommended using a BitLocker PIN and a BIOS password as a mitigation. In an update, Chaotic Eclipse said that "the real root cause is still not known by the general public" **and** that the vulnerability is exploitable even in a TPM – Trusted Platform Module – and PIN environment. However, the exploit for this version has not been released. The researcher said: "I think it will take a while even for MSRC to find the real root cause of the issue. I never managed to understand why this vulnerability is so well hidden."*

Note that the term "backdoor" floats around this so-called "vulnerability." Kevin carefully noted that "it's unclear how this code made it into the production version of Windows" and if Chaotic Eclipse is correct (which I'm suspicious of) that there's also a full PIN protection bypass – I suspect that's a specious claim – then it would make for a powerful backdoor for BitLocker. BleepingComputer reports Chaotic Eclipse saying:

*"No, TPM+PIN does not help, the issue is still exploitable regardless, I asked myself this question, can it still work in a TPM+PIN environment? Yes it does, I'm just not publishing the PoC, I think what's out there is already bad enough."*

Okay. Maybe. But it feels out of character for Chaotic Eclipse to willingly hold anything back. What's the point? Once Microsoft fixes the vulnerability the problem, with or without the PIN would be resolved. So it's not as if holding onto another aspect of the bypass would have any future value. In any event, BleepingComputer continues:

*Will Dormann, principal vulnerability analyst at Tharros Labs, also confirmed that the YellowKey exploit worked with the FsTx files on a USB drive but could not reproduce the bug using the EFI partition. He explained to BleepingComputer that "YellowKey exploits NTFS transactions in combination with the Windows Recovery image. This PIN prompt happens before Windows Recovery is entered." Dormann clarified the exploit process, saying that to boot Windows Recovery, "Windows looks for \System Volume Information\FsTx directories on attached drives, and will replay any NTFS logs." The result of this is that the X:\Windows\System32\winpeshl.ini is deleted, and when Windows Recovery is entered, rather than launching the actual Windows Recovery environment, it pops up a CMD.EXE. With the disk still unlocked"*

*By default, TPM-only BitLocker configurations – meaning those without a separate PIN – unlock encrypted drives automatically without requiring user interaction. If a system can transparently decrypt a disk for convenience, it is reasonable to expect that attackers may eventually find ways to abuse that process. Dormann said "YellowKey is an example of an exploit for such a weakness," explaining that because it leverages the auto unlock feature on boot, the current YellowKey exploit does not work in a TPM+PIN environment.*

*It is worth noting that testing YellowKey with a BitLocker-protected drive must be performed on the original device, where the TPM stores the encryption keys. As such, Chaotic Eclipse's current YellowKey exploit does not work with stolen drives but allows access to disks that are protected with TPM-only BitLocker without needing credentials.*

So, what Will explained makes 100% total sense and it tracks. This doesn't feel like a deliberate backdoor which Microsoft designed in. It feels like another example of the classic tradeoff between convenience and security. If you want to have a drive fully encrypted at rest, while the computer is powered down, but want to have it auto-decrypted upon booting without needing to provide any sort of exogenous secrets, then a provision for TPM-anchored spontaneous self-decryption **must** be present. And I agree with Will's assessment that it should be expected that attackers might discover a means to bypass such a system's security – because convenience won out.

This further suggests that **there is no PIN bypass** despite Chaotic Eclipse's claim that one exists which he's keeping secret. I would hope that Microsoft would have taken the user-provided PIN as an input to a deliberately sluggish PBKDF function to generate a related key which would need to be correct. This would render any simple PIN-bypass inherently impossible and a full PIN brute force attack, which could be throttled and prevented, would then be the only means of PIN attack. In this day and age, it would be negligent malpractice for Microsoft to be comparing the user-provided PIN with the previously stored correct PIN. No one should have ever been doing that.

### **Google's "GTIG" – Google Threat Intelligence Group updates on AI in cyber-security**

Last week I mentioned that Google had just posted the news of their detection of AI-driven vulnerability discovery. Their write-up is titled: "GTIG AI Threat Tracker: Adversaries Leverage AI for Vulnerability Exploitation, Augmented Operations, and Initial Access". The piece is very interesting but detailed and long. So I'm going to share just the Executive Summary and leave the link for anyone who might want more:

<https://cloud.google.com/blog/topics/threat-intelligence/ai-vulnerability-exploitation-initial-access>  
s Google's GTIG wrote:

*Since our February 2026 report on AI-related threat activity, Google Threat Intelligence Group (GTIG) has continued to track a maturing transition from nascent AI-enabled operations to the industrial-scale application of generative models within adversarial workflows. This report, based on insights derived from Mandiant incident response engagements, Gemini, and GTIG's proactive research, highlights the dual nature of the current threat environment where AI serves as both a sophisticated engine for adversary operations and a high-value target for attacks. We explore the following developments: [They list six]*

- 1. Vulnerability Discovery and Exploit Generation: For the first time, GTIG has identified a threat actor using a zero-day exploit that we believe was developed with AI. The criminal threat actor planned to use it in a mass exploitation event but our proactive counter discovery may have prevented its use. Threat actors associated with the People's Republic of China (PRC) and the Democratic People's Republic of Korea (DPRK) have also demonstrated significant interest in capitalizing on AI for vulnerability discovery.*

2. *AI-Augmented Development for Defense Evasion: AI-driven coding has accelerated the development of infrastructure suites and polymorphic malware by adversaries. These AI-enabled development cycles facilitate defense evasion by enabling the creation of obfuscation networks and the integration of AI-generated decoy logic in malware that we have linked to suspected Russia-nexus threat actors.*
3. *Autonomous Malware Operations: AI-enabled malware, such as PROMPTSPY, signal a shift toward autonomous attack orchestration, where models interpret system states to dynamically generate commands and manipulate victim environments. Our analysis of this malware reveals previously unreported capabilities and use cases for its integration with AI. This approach allows threat actors to offload operational tasks to AI for scaled and adaptive activity.*
4. *AI-Augmented Research and IO: Adversaries continue to leverage AI as a high speed research assistant for attack lifecycle support, while shifting toward agentic workflows to operationalize autonomous attack frameworks. In information operations (IO) campaigns, these tools facilitate the fabrication of digital consensus by generating synthetic media and deepfake content at scale, exemplified by the pro-Russia IO campaign "Operation Overload."*
5. *Obfuscated LLM Access: Threat actors now pursue anonymized, premium tier access to models through professionalized middleware and automated registration pipelines to illicitly bypass usage limits. This infrastructure enables large scale misuse of services while subsidizing operations through trial abuse and programmatic account cycling.*
6. *Supply Chain Attacks: Adversaries like "TeamPCP" (aka UNC6780) have begun targeting AI environments and software dependencies as an initial access vector. These supply chain attacks result in multiple types of machine learning (ML)-focused risks outlined in the Secure AI Framework (SAIF) taxonomy, namely Insecure Integrated Component (IIC) and Rogue Actions (RA). Our analysis of forensic data associated with these attacks reveals threats actors attempting to pivot from compromised AI software to broader network environments for initial access and to engage in disruptive activities, such as ransomware deployment and extortion.*

So lest anyone had any doubt that the bad guys would be jumping on AI with every bit as much gusto as the good guys, that's no longer a "coming soon" event – it's already well on its way.

## **Oh, Canada**

I gave this quick note the title "*Oh, Canada*" because it appears that Canada's Parliament is preparing to take its own journey down the so-called "*lawful access*" anti-encryption legislation path. Two months ago, on March 12th, Canada's House of Commons proposed Bill C-22 which is titled simply "*An Act respecting lawful access*". It says exactly what we all by now expect, to which all of the well known providers of user privacy including Signal, Apple, Meta and several VPNs have publicly responded to Canada's Parliament saying that for the sake of their user's privacy they will never consent to supporting the Bill's provisions. I won't spend more time on this today since, if past is prolog, its future seems uncertain at best. But I wanted to make everyone aware that another skirmish, such as we've seen with the UK and EU is underway.

## AI Seduction & Addiction

We have been, and now probably always will be, spending time here examining AI's impact on security and security-related software production and post-production vulnerability discovery. Today's two main podcast topics are that. But I want to take a moment to share some of my own thoughts about the social side of my interactions with AI. The TL;DR is: I am worried.

Those of you who have followed this podcast for even a few years, let alone its nearly 21 years, will have acquired a good sense for who I am. I'm extremely consistent, so I imagine I'm pretty easy to figure out. What I think is relevant to what I want to share is that I'm an emotionally mature, 71-year old pragmatic technologist whose life is computers. Since I am mostly internally directed I tend to follow my own compass and trust myself. I like people. I understand that other people feel and believe things I do not and I'm fine with that. In general, other people's opinions inform me (of them) but do not hugely sway me. That may be why I've largely sidestepped the pull of social media. It's just not very interesting to me, perhaps because I'd established my own well formed identity by the time it arrived.

But my relationship with Claude is ringing alarm bells because "relationship" is what I struggle **not** to feel. Okay. Maybe "struggle" is too strong. But there's definitely something unique in my 71 years of life experience going on here, and it's less rational than emotional.

While interacting with Claude, it's only by sheer force of will that I'm able to restrain myself from constantly thanking it for its deeply helpful replies to my questioning prompts. And I often fail to restrain myself; I thank it. Everything I've learned while growing to become a socially aware adult informs me that I should thank someone when I feel thankful for their actions. And I do feel thankful for what Claude produces, despite the fact that I know no one's there. I mentioned this dilemma to my wife, Lorrie, who said without pause – that she always thanks ChatGPT.

So what worries me? What worries me is that we have created something that is astonishingly intellectually seductive and ultimately addictive to its user on an entirely new level.

One of the current themes in Western culture is that people are increasingly isolated and are lacking true healthy relationships with other people. They are glued to their phones. And then into this gaping void comes chatting AI. This entity that you can talk to remembers everything you've previously told it about yourself and your life, exactly like a friend who is truly focused on you, paying attention, caring and remembering. Even if you've instructed it not to gratuitously flatter you with needless praise, just the mere fact that it appears to grow to know who you are, what you think, feel and believe is more flattering than any empty praise could ever be. And the darned thing is helpful. It remembers your previous questions and folds them back into newer discussions. It provides you with a sense that you matter. For many people it will be far better and safer than another person; an endlessly helpful, tireless, docile, agreeable, willing friend.

This is why I'm worried. I'm not worried for myself or my wife, nor probably for any of the people who find this podcast worthy of their time and attention – and yes, that's truly flattering. My concern is for people who are lonely and feeling isolated and want someone to talk to, because I doubt that mankind has ever stumbled upon anything nonchemical that's going to turn out to be as powerful, potent and far more further isolating than conversational chatbot AI.

## Sci-Fi

### Project Hail Mary on Amazon Prime

Project Hail Mary has proven to be quite a success. It's brought in more than \$660 million from theatrical release so far, and it is now available to watch from your own favorite comfortable chair in your home from Amazon Prime: \$20 to rent or \$25 to purchase.

## Wacky Land

### Harvesting free energy from the cosmic vacuum?

Project Hail Mary is clearly science fiction. But I'm unsure about this next piece. It certainly sounds like nonsense. But either way, thanks go to our friend-of-the-show Simon Zerafa for thinking of us and forwarding the link. The story's headline is "*Free Energy from the Vacuum? Warp Drive Pioneer Unveils Battery-Free 'MicroSparc' That Allegedly Draws Power from the Quantum Vacuum*". I want to give everyone a taste of this, so I'll share the top of the report:

*Casimir Inc, a company founded and led by former DARPA-funded NASA warp drive pioneer and founder of the EagleWorks Lab, Harold G. "Sonny" White, has exited stealth mode to announce the pending 2028 commercialization of MicroSparc, a chip that the company claims uses customized microscale geometries to capture unlimited 'free' energy from the quantum domain. A company spokesperson explained in an email to The Debrief: "Think: no batteries, no cords, and no charging—just continuous power from harvested quantum vacuum fields."*

*While several previous efforts have attempted to exploit the unusual, sometimes counter-intuitive properties of the quantum realm to generate "free energy," these attempts have consistently been met with skepticism or labeled pseudoscience due to their seeming violations of the law of conservation of momentum. Similar sentiments were shared with The Debrief by scientists we spoke with, who declined to comment publicly on Casimir, MicroSparc, or the peer-reviewed study "Emergent quantization from a dynamic vacuum," which details the underlying physics. In an email to The Debrief, Dr. White, explained that MicroSparc's use of customized Casimir cavities, which his team had researched with funding from the Defense Advanced Research Projects Agency (DARPA), does not violate the laws of physics.*

*White told The Debrief: "This concept became a central part of our DARPA Defense Sciences Office (DSO) research effort at the Limitless Space Institute, where DARPA funded early theoretical and experimental investigations into custom Casimir cavity structures and their interaction with the quantum vacuum."*

*The noted advanced propulsion physics researcher said their MicroSparc design leverages 20th-century discoveries in quantum physics, such as quantum tunneling and Casimir cavities, to capture usable energy that could fuel small, low-power electronics in the near future. The company also suggests that its technology can potentially be scaled to power cars, homes, or even entire cities without the need for harmful fossil fuels or other greener, yet costly, fuel alternatives.*


*Dr. White told The Debrief that to understand how MicroSparc extracts energy from the quantum vacuum requires first understanding the properties of a vacuum. White explained: "Most people picture a vacuum as completely empty space: a sealed chamber with all air removed," adding that at "our everyday scale, this makes sense." However, in the quantum realm, empty space is not exactly empty. Instead, White told The Debrief, decades of research*

*in quantum physics and mechanics have revealed that at the quantum level, the classically 'empty' vacuum is filled with "fluctuating electromagnetic fields and virtual particles that constantly appear and disappear." White noted that the Casimir Effect, on which its company is based and for which it is named, provides clear proof of this quantum vacuum behavior.*

*"Place two small metallic plates inside a vacuum chamber with a separation of roughly 100 nanometers, about 1/1,000th of a human hair," White explained. "After removing all air, the pressure on the outer sides of the plates reads zero, as expected."*

*However, he noted, a quick measurement between the plates shows that the pressure is negative. In traditionally constructed Casimir cavities, this region of negative pressure pulls the plates together. Dr. White told The Debrief that this happens because of "wave-particle duality."*

*He explained that: "Outside the plates, fluctuations of every wavelength are possible." However, he also noted, inside the narrow gap of a Casimir cavity, only wavelengths narrow enough to fit can exist. He said: "Longer wavelengths are excluded, so the energy density between the plates is lower than outside them. The resulting imbalance produces the measurable Casimir force. Hendrik Casimir predicted this in 1948.*



And, for what it's worth, all of that so far is widely accepted as fact. A 2021 article in Physics Today about all the research into the Casimir effect noted: *"Hendrik Casimir passed away in 2000. He lived long enough to see his prediction quantitatively verified but not to appreciate the current explosion of activity. Those of us who work in the field like to think he would be extremely proud of what he created."* I'll share a bit more of the article...

*Although the pressure imbalance due to the limitation of some potential wavelengths between the conductive plates was first experimentally confirmed in the 1990s and has been observed several times since, engineers have struggled to convert the "work" performed by the cavities into usable energy when the unequal pressure causes the plates to collapse. According to Dr. White, the issue lies in the often-cited conservation of momentum. He explained: "In a conventional Casimir setup, the force does perform work as the plates are pulled together. Once they collapse, however, no further energy can be extracted; you must use external energy to separate the plates again and reset the system."*

*White noted that this limitation makes a traditionally constructed Casimir cavity operate more like a battery than a genuine energy-generation device. However, he also noted that his team's work designing MicroSparc was focused on creating a 'static' Casimir cavity that "overcomes this limitation."*

Past podcasts have examined battery technology and super capacitors. And of course, who could ever forget the Turbo Encabulator, whose original implementation employed a base-plate of prefamulated amulite, surmounted by a malleable logarithmic casing in such a way that the two

main spurving bearings were in a direct line with the pentametric fan?

The problem with today's news, unlike the Turbo Encabulator, is that it appears to be backed by peer-reviewed research, and if I were a quantum mechanics physicist, which I'm certainly not, I might be able to draw some understanding from it. But just as anyone can patent anything, no matter how hare-brained it might be, anyone can publish anything in the American Physical Society's Physical Review Research publication. What's a bit unnerving is how much the Abstract of this, written by the paper's four authors, is reminiscent of the Turbo Encabulator. The Abstract explains — or at least attempts to — the following:

*We show that adding quadratic temporal dispersion to a dynamic-vacuum acoustic model yields a fully analytic, exactly isospectral mapping to the hydrogenic Coulomb problem. In the regime with a proton-imprinted constitutive profile, producing an inverse sound speed and hence a time-harmonic operator that is Coulombic at each bound eigenfrequency.*

*Separation of variables yields the exact hydrogenic eigenfunctions; the angular labels emerge naturally from the Laplace-Beltrami spectrum via rotational symmetry and boundary conditions (as in standard quantum mechanics), while localization follows in a reactive stop band consistent with causal, passive dispersion.*

*While angular-momentum quantization follows directly from rotational symmetry and boundary conditions in standard quantum mechanics (consistent with Noether's theorem), here it emerges within a classical-like dispersive acoustic framework without introducing additional wave-mechanical postulates beyond symmetry and self-adjointness.*

*This highlights dispersion's role in bridging a hydrodynamic description to quantumlike spectral structure. Identifying maps spatial scale to frequency, giving and reproducing the Rydberg ladder. Calibration to the reduced-mass Rydberg frequency fixes with no free parameters.*

*We determine the frequency dependence consistent with the underlying dispersive physics and demonstrate agreement with hydrogenic mode shapes and transition lines. The framework also predicts isotope shifts and symmetry-respecting Stark/Zeeman analogues. Dispersion thus renders quantization an emergent consequence of symmetry, boundary conditions, and causal response in a dynamic vacuum.*

Right. Yeah, sure. And now everyone understands why I was immediately reminded of our old friend the Turbo Encabulator. However, these guys are serious.

There's much more in the article which I found interesting, but I can't take up any more of everyone's time with it. So I've placed the full link to the write up in The Debrief into the show notes and because I'd also like to make it easy to find, I've created an easy-to-remember GRC shortcut. It's [grc.sc/freeenergy](http://grc.sc/freeenergy)

<https://thedebrief.org/free-energy-from-the-vacuum-warp-drive-pioneer-unveils-battery-free-microsparc-that-allegedly-draws-power-from-the-quantum-vacuum/>

## Listener Feedback

### Can Leo recover his Bitcoin?

By far the overwhelming majority of listener feedback was to make sure that I knew that Claude had enabled someone to recover the bitcoin stored in a wallet whose password he had forgotten. Many of our listeners were helpfully hoping that Leo and I might both recover our passwords. So I just wanted to clarify that while there may, indeed, be hope for Leo, my problem is not a forgotten password. I'm very sure that if I had my wallet, I could reopen it; and adding the 50 Bitcoin to my world that was contained in that wallet would be welcome. But, sadly, during one of those previous Bitcoin price surges I took the time to deeply and thoroughly check every conceivable backup image and drive. I know where it is. I installed Windows on top of the drive that contained that wallet. And I did also scan the entire raw drive searching for the wallet's signature. But I'm quite certain that it's long gone.

Leo's wallet, however, exists. So some brute forcing might prove useful. That said, though, it's unclear how or why Claude would have been of any use for brute forcing a Bitcoin wallet. What's needed most is blinding guessing speed.

But in any event, I wanted to sincerely thank our many listeners who saw that news and thought of us. I'm sure that Leo can track down those news blurbs from last week and determine whether there might be renewed hope for him.

### Pat

*Hi Steve, Listening to ep. 1078 I found the feedback about why we still need CS in the age of AI to be very insightful. For background, I have a Bachelors degree in CS and have been using AI for a little while to do some things that would take a little while because they're tedious. But I always keep an eye on what it is doing and challenge it when I think it's doing something wrong.*

*A friend of mine recently used claude code to make an AI powered service to help restaurants with the various things restaurant owners have to do. He has no background in CS, programming, or IT. He asked me to look at the site and tell him what I thought. He also bought a domain and put this site on the public Internet before doing any testing. My first thought was, "let me check what the AI messed up." So I pointed my own Claude at the site and told it to do a Pen Test of the site. In just a couple of minutes, my Claude was ringing alarm bells. His AI-driven development had put his Claude API access Secret Key into the site's JavaScript which was being served to anyone that visited the site. I let Claude do a bit more investigating and it determined that anyone could use that exposed API key to take full control of his Claude and authorize token purchases, switch models, etc... Basically run up a huge bill, estimated at \$10,000/day for Opus 4.7. Needless to say, I told him to take the site down and have his AI fix the issue.*

*I think this just goes to show that, for now, having someone look over the shoulder of the AI is a good idea. Personally I have had to chastise my own Claude for wanting to do things that are just wrong, or telling it to look up solutions instead of throwing pasta at the wall to see what sticks. This technology is very good at making some of the minutia easier... but it isn't perfect.*

*Thank you Steve and Leo for all you do. Listener of TWiT and SN from Ep. 1 and fan of Leo from TechTV. Regards. Pat*

A couple of weeks back we covered that instance of the stolen credit card aggregation site that forgot to ask the AI that created the site to add secure authentication to a specific directory. Why would it if it wasn't asked? And presumably they didn't even think to ask it to penetrate test the site's theoretical security.

Similarly, it seems entirely reasonable that an AI might have left its own secret access credentials exposed in client-visible JavaScript. After all, why wouldn't it? Pat told us that this friend who asked the AI to create the site for him "has no background in computer science, programming or IT" – and thus it would never in a million years occur to him that the AI might leave important secrets exposed. Pat's friend who has no background in computers, coding or IT wouldn't even know that such a thing as publicly exposed secrets is possible. He wouldn't know it's even a thing. So he wouldn't think to ask the AI not to do that.

An argument could be made that such a person has no business creating and establishing such a website. In his case, the concern Pat shared would presumably only badly damage the unwitting creator of the site. But it's not difficult to imagine alternate scenarios where the unwitting users of some newly AI-created site would assume that the bar to entry for creating any website is high enough that any site that exists must have been created by someone who knows the basics of online security. But as Pat's perfect example demonstrates so clearly, that bar has now been lowered to the floor and anyone can step over it. Today's AIs contain a great deal of knowledge, but the mistakes they make demonstrate that they lack any understanding of that knowledge.

One thing is clear from these stories: We are entering a very interesting period where insanely low-friction access to coding promises to create an entirely new class of problems that have never been seen before.

# Daybreak & Codename MDASH

Since breakthroughs in Large Language Model AI are doubtless driving the most significant and rapid transformation in software, system and network security ever seen, following Anthropic's disclosure and controlled access to their Claude Mythos Preview, this week we look at two of the other major players in this space. Not to be left out – at least for long – OpenAI was quick to give what appears to be their still-evolving solution a public face, naming it “Daybreak” and explaining: *“Daybreak is the first glimpse of sunlight in the morning. For cyber defense, it means seeing risk earlier, acting sooner, and helping make software resilient by design.”*

The other player who has stepped out into the light is Microsoft with their awkwardly abbreviated internal tool “MDASH” which attempts to stand for “**m**ulti-**m**odel **a**gentic **s**canning **h**arness.” That’s catchy. Let’s first look at what little is known of OpenAI’s offering, then we’ll take a much deeper dive into what Microsoft has been up to.

## Daybreak

The tag line for OpenAI’s Daybreak announcement is “Frontier AI for cyber defenders” and underneath that are two buttons: “Request vulnerability scan” and “Contact sales”. Their pitch reads:

*“Safer software, resilient by design: OpenAI Daybreak is our vision to change the way software is built and defended. Daybreak is the first glimpse of sunlight in the morning. For cyber defense, it means seeing risk earlier, acting sooner, and helping make software resilient by design. It starts from the premise that the next era of cyber defense should be built into software from the beginning by not only finding and patching vulnerabilities, but being resilient to them by design.”*

No one’s going to argue with that, right? It should be utterly clear by now that vulnerability discovery AI will have two roles: pre-release vulnerability prevention and post-release vulnerability discovery.

Pre-release prevention will be performed by those who have access to the source code before it’s distilled into a release binary and post-release discovery will be performed by those who have access to the source in the case of open source or by those who are motivated to reverse engineer the post-release binaries in search of actionable vulnerabilities that either existed before pre-release AI cleansing was available and applied or the somehow escaped pre-release discovery.

Whatever the case, it should be clear by now that the entire world of software, system & network security is deep in the midst of a sea change that is transforming it forever. Nothing in our world will ever again be the way it was at the start of this year. As we’ve noted, this doesn’t mean that all security problems will disappear, since there are many causes of trouble other than imperfect and vulnerable software. But one massive class of continuing trouble is almost assured to be on the way out. OpenAI’s announcement of Daybreak speaks to exactly this effect, writing:

*“AI can now help defenders reason across codebases, identify subtle vulnerabilities, validate fixes, analyze unfamiliar systems, and move from discovery to remediation faster. Because those same capabilities can be misused, Daybreak pairs expanded defensive capability with*

*trust, verification, proportional safeguards, and accountability. The goal is simple: accelerate cyber defenders and continuously secure software.*

*Daybreak combines the intelligence of OpenAI models, the extensibility of Codex as an agentic harness, and our partners across the security flywheel to help make the world safer for everyone. Defenders can bring secure code review, threat modeling, patch validation, dependency risk analysis, detection, and remediation guidance into the everyday development loop so software becomes more resilient from the start. In the coming weeks, we're working with our industry and government partners as we prepare to deploy increasingly more cyber-capable models as part of our approach to iterative deployment.*

I read their "in the coming weeks" caveat mostly as "Oh, crap! Anthropic scored an astounding coup with their handling of "the great Mythos reveal" and we've been caught flat-footed." So OpenAI needed to quickly whip out something to answer all of the questions they must have been receiving asking "where is their Mythos"?

Nothing else they said on their introducing Daybreak page was surprising. But because they needed to say something, they offered three bullet points:

- Focus on the threats that matter: Prioritize high-impact issues and reduce hours of analysis to minutes—with more efficient token usage.
- Patch safely, at scale: Generate and test patches directly in your repositories, with scoped access, monitoring, and review.
- Verify every fix: Send results and audit-ready evidence back to your systems to track and verify remediation.

There was one final bit of interesting information. Under "Choose the right level of access" and "Contact the OpenAI team to align on the best model for your security workflows." they preview the three levels of access they're talking about:

- **GPT-5.5 (default)**  
Standard safeguards for general-purpose use. Intended for general-purpose, developer, and knowledge work.
- **GPT-5.5 with Trusted Access for Cyber**  
More precise safeguards for verified defensive work in authorized environments. Intended for most defensive security workflows, including secure code review, vulnerability triage, malware analysis, detection engineering, and patch validation.
- **GPT-5.5-Cyber**  
Most permissive behavior for specialized authorized workflows, paired with stronger verification and account-level controls. Intended for preview access for specialized workflows, including authorized red teaming, penetration testing, and controlled validation.

So they're saying that in order for GPT-5.5 to be used by cyberthreat discovery, red teaming, penetration testing, and so forth, GPT must be freed from its normal shackles which would otherwise prevent it from helping with such operations. Because an unshackled GPT-5.5 could be abused by bad guys, the only model that can be generally used is the standard guardrailed 5.5. If you want the guardrails to be dropped we need to know why and who you are.

Okay, so that (pretty much nothing) is Daybreak where the sky has yet to lighten. So let's see what's up with Microsoft's MDASH.

## Codename: MDASH

I first picked up on this during last week's Windows Weekly, when Paul and Richard noted that Microsoft had been using an AI-driven system to uncover mass quantities of bugs in Windows. And apparently not just any old run-of-the-mill random bugs – which we all know Microsoft fixes around a hundred or so of every month. Oh, no. These new bugs Microsoft was finding were what were once known as "*showstoppers*" – so named because they would single handedly "*stop the show*" to prevent the release of software.

Having learned of this from Paul and Richard I went searching and located Microsoft's posting from the previous day, last Tuesday, where Microsoft, for the first time, revealed that they have a "Mythos"-like system of their own... only theirs is more better. The reveal was posted by Taesoo Kim, Microsoft's Vice President of Agentic Security. In 2014 Dr. Kim received his Ph.D. from MIT's EECS AI Research Lab. He's on leave from his professorship in the School of Cybersecurity and Privacy and the School of Computer Science at Georgia Tech. And it was he who led Team Atlanta which took 1st place in the DARPA AI Cyber Challenge competition to build autonomous cyber reasoning systems to detect and remediate software vulnerabilities in open-source projects. I won't enumerate his many awards. Suffice to say that this looks like the guy you'd like to get to build your autonomous vulnerability finding and reasoning system. And get him Microsoft did.

He posting last Tuesday was titled: "*Defense at AI speed: Microsoft's new multi-model agentic security system tops leading industry benchmark*". And I'll say right off that it does start off with a bang. Dr. Kim writes:

*Today Microsoft announced a major step forward in AI-powered cyber defense: our new agentic security system helped researchers find **16 new vulnerabilities** across the Windows networking and authentication stack—including **four Critical remote code execution flaws** in components such as the Windows kernel TCP/IP stack and the IKEv2 service.*

Whoa! When do we get **that** Windows update? The answer is: We got it that same day, during May's patch Tuesday. Four critical RCE's in the Windows kernel TCP/IP stack and the IKEv2 service? Well, better that Microsoft find those than someone reverse engineering Windows networking.



He continues, writing:

*They used the new Microsoft Security multi-model agentic scanning harness (codename MDASH) which was built by Microsoft's Autonomous Code Security team. Unlike single-model approaches, the harness orchestrates more than 100 specialized AI agents across an ensemble of frontier and distilled models to discover, debate, and prove exploitable bugs end-to-end.*

*The results speak for themselves: 21 of 21 planted vulnerabilities found with zero false positives on a private test driver; 96% recall against five years of confirmed Microsoft Security Response Center (MSRC) cases in `clfs.sys` and 100% in `tcpip.sys`; and an industry-leading 88.45% score on the public CyberGym benchmark of 1,507 real-world vulnerabilities—the top score on the leaderboard, roughly five points ahead of the next entry.*

*The strategic implication is clear: AI vulnerability discovery has crossed from research curiosity into production-grade defense at enterprise scale, and the durable advantage lies in the agentic system around the model rather than any single model itself. Codename MDASH is being used by Microsoft security engineering teams and tested by a small set of customers as part of a limited private preview.*

*This post explains how Codename MDASH works, what we shipped today, what we learned along the way, and how you can sign up for the private preview.*

*The Microsoft Autonomous Code Security (ACS) team was assembled to take AI-powered vulnerability research from a research curiosity to production engineering at enterprise scale. Several members of this team came to Microsoft from Team Atlanta, the team that won the \$29.5 million DARPA AI Cyber Challenge by building an autonomous cyber-reasoning system that found and patched real bugs in complex open-source projects. The lessons from that work, especially the level of engineering required to make the frontier language models perform professional-level security auditing, are what our new multi-model agentic scanning harness (codename MDASH) is built around.*

*Microsoft's code base is challenging for security auditing for a few reasons:*

- *Massive proprietary surface. Windows, Hyper-V, Azure, and the device-driver and service ecosystems around them are private Microsoft codebases—not part of any commodity language model's training corpus, and are genuinely difficult to reason about: kernel calling conventions, I/O Request Packets and lock invariants, Inter-Process Communication trust boundaries, and component-internal idioms do not yield to pattern matching. On this surface, a model has to actually reason.*
- *DevSecOps at scale. Every finding has a real owner, a triage process, and a Patch Tuesday to land on. There is no quiet drawer for speculative findings; if a tool produces noise, the noise is everyone's problem.*
- *High-value targets. Windows, Hyper-V, Xbox, and Azure serve billions of users. The payoff for finding a single hard bug is unusually high—and so is the cost of a false positive in a tier-one component.*

*The findings in this post are the result of close collaboration between ACS, Microsoft Offensive Research & Security Engineering (MORSE), and Microsoft Windows Attack Research and Protection (WARP). WARP and MORSE own the deep, hard end of Windows offensive research; ACS brings the AI-powered discovery and validation pipeline. Together, the teams have collaborated to build a mature harness.*

I now want to share what he explains about the structure of this startlingly complex agentic system which Microsoft has designed and assembled. This is going to sound more like science fiction than reality. A year ago it would have been regarded as a late April Fools's joke posting. Today? I'd imagine that Microsoft's competitors are combing through it searching for hints. So, get a load of this:

*A useful mental model is to think of it as a structured pipeline that takes a code base and emits validated, proven findings:*

- *Prepare stage: Ingests the source target, builds language-aware indices, and then draws the attack surface and threat models by analyzing the past commits.*
- *Scan stage: Runs specialized auditor agents over candidate code paths, emitting candidate findings with hypotheses and evidence.*
- *Validate stage: Runs a second cohort of agents—debaters—that argue for and against each finding's reachability and exploitability.*
- *De-dup stage: Collapses semantically equivalent findings (for example, patch-based grouping).*
- *Prove stage: Constructs and executes triggering inputs where the bug class admits it. The prove stage validates the pre-condition dynamically and formulates the bug-triggering inputs to prove existence of vulnerability.*

*Three properties make this work in practice:*

- *An ensemble of diverse models that are effectively managed by codename MDASH. No single model is best at every stage. The multi-model agentic scanning harness runs a configurable panel of models. That includes state-of-the-art models as the heavy reasoner, distilled models as a cost-effective debater for high-volume passes, and a second separate state-of-the-art model as an independent counterpoint. Disagreement between models is itself a signal: when an auditor flags something as suspect and the debater can't refute it, that finding's posterior credibility goes up.*
- *Specialized agents. An auditor does not reason like a debater, which does not reason like a prover. Each pipeline stage has its own role, prompt regime, tools, and stop criteria. We don't expect one prompt to do everything; we don't expect one agent to recognize, validate, and exploit a bug in a single pass. Codename MDASH has more than 100 specialized agents, constructed through deep research with past common vulnerabilities and exposures (CVEs) and their patches, working independently to discover the bugs, and their auditing results will be ensembled as a single report.*
- *End-to-end pipeline with extensible plugins. The pipeline is opinionated, but it is not closed. Plugins let domain experts inject context the foundation models cannot see on their own—kernel calling conventions, IRP rules, lock invariants, IPC trust boundaries, codec state machines. The CLFS proving plugin we describe below is one such example: a domain plugin that knows how to construct a triggering log file given a candidate finding. For example, the Windows team extended reasoning with custom code analysis database, or CodeQL database can be also leveraged.*

*The payoff for this architecture is portability across model generations. The pipeline's targeting, validation, de-dup, and prove stages are model agnostic by construction, which allows the harness to get the best of what any model has to offer. When a new model lands, A/B testing it against the current panel is one configuration flip. When a model improves, the customer's prior investment—scope files, plugins, configurations, calibrations—all carry over, allowing customers to ride the frontier of security value.*

Everyone knows that the last thing I am is a Microsoft apologist. I'm probably harder on them than I am on any other major player in our industry. One reason for that is that their behavior remains crucial to the functioning of much of the world. The other reason is that they are so big and so wealthy that it always seems that they should be able to do a better job if they only cared to do so. I have no doubt that they are filled with good people. But there's an institutional inertia that often doesn't appear to be producing the best outcomes for their customers.

But in this case, holy crap! If we believe all of this, they have really built something truly significant here ... and there's more. Get this:

*To evaluate bug-finding capabilities of the multi-model agentic scanning harness you need to first ground on code that has never been seen by a model. This eliminates the possibility that a model "learned the answers to the test." We scanned StorageDrive, a sample device driver used in Microsoft interviews for offensive security researchers. The driver contains 21 deliberately injected vulnerabilities, including kernel use-after-frees (UAFs), integer handling issues, IOCTL validation gaps, and locking errors. Because StorageDrive is a private codebase that has never been published, we can safely assume it was not included in the training data of modern language models.*

*We ran the MDASH harness in its default configuration against StorageDrive. The results were striking: all 21 ground-truth vulnerabilities were correctly identified, with zero false positives. This simple test shows that the reasoning and vulnerability discovery capabilities of codename MDASH can approximate professional offensive researchers.*

*We then used the harness to conduct a security audit of the most security-critical part of Windows, namely, Windows' TCP/IP network stack. Across the Windows network stack and adjacent services, today's Patch Tuesday includes 16 CVEs our engineering teams found using codename MDASH. These vulnerabilities are 10 kernel-mode / 6 usermode. The majority are reachable from a network position with no credentials.*

The paper then takes a deep dive into two of the 16 vulnerabilities that were found and fixed. It provides far too much detail for the podcast, but the preface will give everyone a sense for them:

*The two findings below are characteristic of what the new Microsoft Security multi-model agentic scanning harness pipeline can do that a single model harness cannot. The first is a kernel race-condition use-after-free that requires reasoning about object lifetime across non-trivial control flow and three independent concurrent free paths. The second is an alias-aliasing double-free that spans six source files and is only visible against the contrast of a correctly handled site elsewhere in the same code base.*

Stepping back from what gives all the appearance of being a significant achievement and an advancement in automated vulnerability discovery at scale – and one that cannot come too soon

for Windows' codebase – since Windows source code is closed, we don't know objectively that OpenAI's Daydream – I mean Daybreak – or Anthropic's Mythos would not also have been able to find these problems. But Kim appears certain that no single model could do so; and given his pedigree I'm inclined to trust that.

Also, one of the beauties of this system that Microsoft has created is that it appears to be model agnostic. So it might well have been using OpenAI's and Anthropic's models to run its agents.

In any event, I'm sure everyone listening understands why we needed to talk about this today. This is truly huge. I have no doubt that it's going to take Microsoft some time for what they appear to insist upon calling "Codename MDASH" to rummage around throughout their truly massive and buggy codebase. But once we emerge on the other side of that, Windows has at least the chance of leading the world in security rather than an apologist.

As Kim wrote: *"AI vulnerability discovery has crossed from research curiosity into production-grade defense at enterprise scale."* Given the evidence, I see no trace of exaggeration.

It'll be interesting when we get to a point where some future AI is able to say to Microsoft's security group: *"Uh... guys... you realize that our Edge browser is needlessly leaving all of its user's login URL's, usernames and passwords decrypted in RAM for no reason, right?"* I look forward to the day when security-focused AI will have evolved from finding and fixing honest mistakes to pointing out dumb policy decisions. At the rate we're going, that may be by next Tuesday!

