# Security Now! #1018 - 03-25-25

## The Quantum Threat

### This week on Security Now!

The dangers of doing things you don't understand. Espressif responds to the claims of an ESP32 backdoor. A widely leveraged mistake Microsoft stubbornly refuses to correct. A disturbingly simple remote takeover of Apache Tomcat servers. A 10/10 vulnerability affecting some ASUS, ASRock and HPE motherboards. Google snapped up another cloud security firm but paid a price! RCS messaging to soon get full end-to-end encryption (done right!). How did an AI Crypto Chatbot lose $105,000? ... and what is an AI Crypto Chatbot? Looks like Oracle may take stewardship of TikTok to keep it in-country. Whoops! 23andMe is sinking — don't let them take your genetics with them! The White House says "the cyber guys should stay!" AI project failure rates are on the rise. Anyone surprised? We then have some relevant listener feedback, and a very interesting update on just how looming is the threat from quantum computing?

## Once seen, never forgotten:

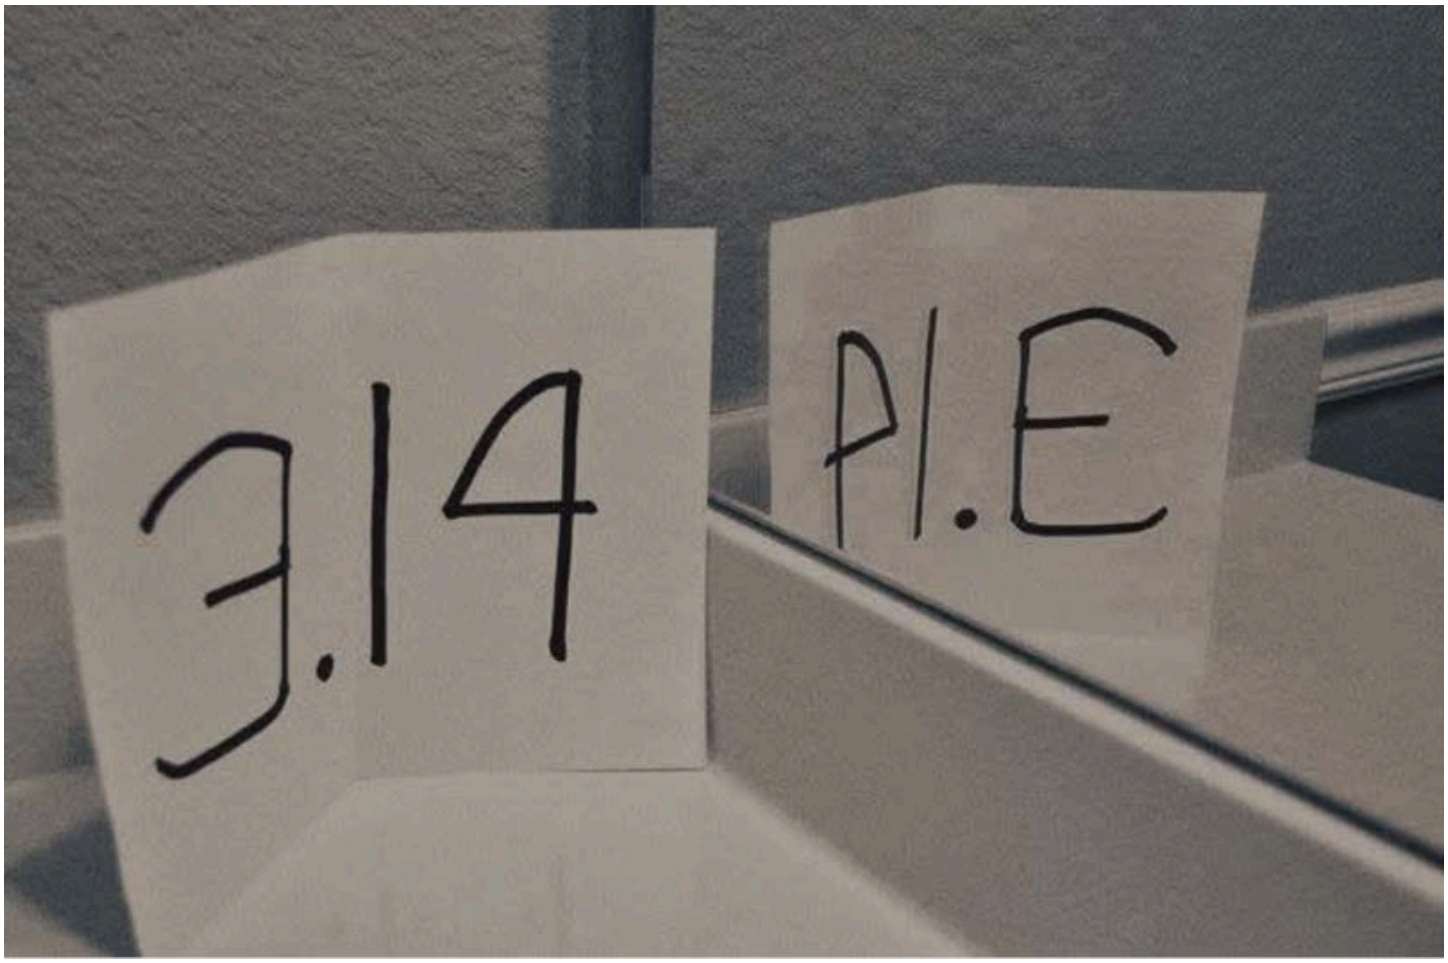

# Security News

**Don't try this at home (or anywhere else, for that matter!)**
Over 100 auto dealerships were being abused in a supply chain attack from a compromised shared video service that is unique to dealerships. When active, the attack presented dealership visitors with a ClickFix webpage which led to a SectopRAT malware infection.

When the user visited any of the over 100 dealerships, there was a chance that a specific JavaScript would load, containing malicious code. If it did, it redirected the user to a page on a compromised host that prompted the user with the increasingly familiar reCAPTCHA claim of "I'm not a robot":

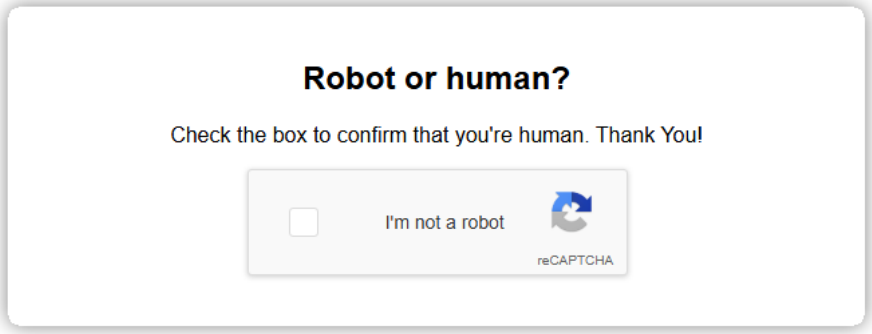


Clicking the "I'm not a robot" checkbox actually gave the lurking JavaScript the privilege to copy a malicious execution string onto the system's clipboard. The malicious Javascript then also extended the "I'm not a robot" dialog with a drop-down flap containing additional verification steps for the user to perform:

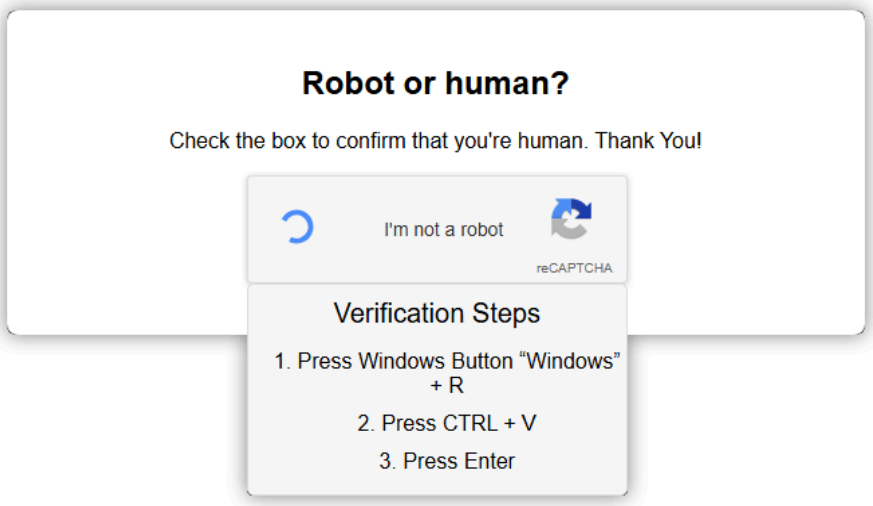


The visitor was instructed to open the Windows RUN dialog by pressing the keyboard's Windows key then 'R'. Doing that would give the RUN field the system's focus. The next instruction was to press CTRL+V which would, of course, copy the malicious command that the JavaScript had pre-loaded onto the system's clipboard into the Run command field. And, finally, the user was instructed to press Enter.

If the user performed these steps, a Powershell script was executed on the user's machine that downloaded further payloads and ultimately installed the remote access trojan SectopRAT.

I know I've mentioned this hack before. But I'm deliberately revisiting this, both because it's so diabolically clever, and because I believe it perfectly captures a significant and fundamental problem that doesn't have any simple solution – and that's the human factor.

I doubt that the listeners of **this** podcast would blindly follow these instructions. But we would all at least pause to consider what's going on here. The important point is that tech-savvy PC users are in the clear minority. As I've noted many times, we're not the target of these social engineering attacks. The vast majority of PC users really have no idea what's going on at all. ***"Instruction following"*** has always been their way of life within the PC world. They may be a brain surgeon by education, training and experience, but that would not prepare them for all of the many clever ways a PC user can be tricked into doing something self-destructive.

The great annoyance for me is that I cannot see a future where this is resolved. The only thing I can see which might resolve this – and I'm actually not kidding – would be an entirely different user interface experience for our PCs, where active AI agents interface the user to their personal computation and communications devices. That may sound far-fetched, but I think it's not so much. Once upon a time, all interaction with computers was via a teletype which had a keyboard and typed text from a wide roll of paper. A big jump was to the textual video display screen which was faster and so much quieter. For many years that's all we had. That's all there was. The next big change was to a graphical display which we interfaced to not only with the same textual keyboard but also with the game-changing mouse and on-screen pointer.

My point is that entirely new user-interface paradigm changes have happened in the past. So it's not as if it's not possible for them to happen again. And we already appear to be creeping forward into this next generation. We know what a powerful impediment to change inertia creates. But it does seem clear that future computation and communication workstations won't have a Windows+R command to get users into trouble.

**Espressif responds to the Spanish researcher's "backdoor" discovery:**

*Shanghai, China: Recently, some media have reported on a press release initially calling out ESP32 chips for having a "backdoor". Espressif would like to take this opportunity to clarify this matter for our users and partners.*

*Recently, some media have reported on a press release initially calling out ESP32 chips for having a "backdoor". Of note is that the original press release by the Tarlogic research team was factually corrected to remove the "backdoor" designation. However, not all media coverage has been amended to reflect this change.*

***What was found:***

*The functionality found are debug commands included for testing purposes. These debug commands are part of Espressif's implementation of the HCI (Host Controller Interface) protocol used in Bluetooth technology. This protocol is used internally in a product to communicate between Bluetooth layers. Please read our technical blog to learn more.*

***Key clarification points***

- *Internal Debug Commands: These commands are meant for use by developers and are not accessible remotely. Having such private commands is not an uncommon practice.*
- *No Remote Access: They cannot be triggered by Bluetooth, radio signals, or over the Internet, meaning they do not pose a risk of remote compromise of ESP32 devices.*

- *Security Impact: While these debug commands exist, they cannot, by themselves, pose a security risk to ESP32 chips. Espressif will still provide a software fix to remove these undocumented commands.*
- *Scope: If ESP32 is used in a standalone application and not connected to a host chip that runs a BLE host, the aforementioned HCI commands are not exposed and there is no security threat.*
- *Affected Chipsets: These commands are present in the ESP32 chips only and are not present in any of the ESP32-C, ESP32-S, and ESP32-H series of chips.*

***Espressif's commitment***

*Espressif has always prioritized security and is actively working on continuous product security improvements. We have a standard Product Security Incident Response Process with underlying bug bounty program that is active since 2017. This program offers a bug bounty, encouraging researchers to collaborate with us to discover and fix potential issues, enhancing the security of the entire ecosystem.*

*Espressif also extends its gratitude to the security research community for promptly clarifying that the disclosure does not constitute a backdoor. Their responsible disclosures and continued support have been invaluable in helping users accurately assess the security implications and maintain the integrity of their connected devices.*

*At the same time, we recommend that users rely on official firmware and regularly update it to ensure their products receive the latest security patches. Should you have any questions, please feel free to contact Espressif's official support channels.*

So this is exactly what we concluded from an examination of the location and nature of these so-called "backdoor" commands. The key is that they were never externally accessible. They were simply commands for the internal native Bluetooth HCI controller.

So that wasn't a problem. Here's something that is:

**Eleven Advanced Persistent Threat groups are known to be abusing a Windows 0-day** but because what they are doing is not technically leveraging a Windows flaw, so far, although this was reported to Microsoft by Trend Micro's ZDI – Zero-Day Initiative – six months ago, last September, Microsoft has declined to patch the issue.

We talked about this at the time because it was just a head shaker that in 2024, let alone still today in 2025, Windows LNK link files are still being exploited. And what's more, despite the fact that the exploitation of this single 0-day vulnerability goes back eight years, Microsoft says "no fixie." The eleven APT groups operate out of North Korea, Iran, Russia, and China – none who have recently been behaving as friends of the West. They have all used this 0-day to hide their malicious instructions in LNK files sent to targets and Trend Micro has discovered nearly 1,000 malicious LNK files abusing the technique. Microsoft's response is that it's all working just the way they want it to.

As I said, we covered this before. Recall that there was (and, unfortunately, still is) a way to format the Fields of the LNK file to essentially "white space pad" the actual content of the LNK field so far to the right that none of it shows up when the user goes to examine the LNK file's properties. So the user won't see that they're going to run "EvilMalwareDownloader" when they click the link.

https://www.trendmicro.com/en_us/research/25/c/windows-shortcut-zero-day-exploit.html

I have the link to Trend Micro's fully detailed report in the show notes for anyone who's interested. The high-priority takeaway for our listeners is to NEVER click any link that has an apparently empty Target field – because the Target cannot really be empty. That field must be non-empty for the link to have any effect. So it makes no sense for Target to ever be blank. Never make the mistake of assuming that a blank field means the entire link is benign just because it's not obviously nefarious. It might just as easily be heavily space-padded.

**Apache Tomcat servers are in trouble again.**
The API security firm Wallarm posted an announcement last week titled *"One PUT Request to Own Tomcat: CVE-2025-24813 RCE is in the Wild"*. They wrote:

> *A devastating new remote code execution (RCE) vulnerability, CVE-2025-24813, is now actively exploited in the wild. Attackers need just one PUT API request to take over vulnerable Apache Tomcat servers. The exploit, originally published by a Chinese forum user iSee857, is already available online.*

So here's what we know: This newly disclosed attack leverages Tomcat's default session persistence mechanism along with its support for partial PUT requests. Tomcat is Apache's Java web application server that provides a "pure Java" HTTP web server environment in which Java code can run. This new exploit works within this environment, requiring just two simple steps.

First, the attacker starts by sending a PUT request to upload a malicious session file to the server. The payload is a base64-encoded ysoserial gadget chain that's designed to trigger remote code execution when it's deserialized. This initial PUT request writes a file inside Tomcat's session storage directory. Because Tomcat automatically saves session data in files, the malicious payload is now stored on disk, waiting to be deserialized. So the first step essentially causes the Apache Tomcat server to upload and store the attacker's Java attack file.

Then, with the session file is uploaded, the attacker triggers deserialization of that file by sending a simple GET request by providing a JSESSIONID cookie which points to the malicious session:

```
GET / HTTP/1.1
Host: vulnerable.host:8080
Cookie: JSESSIONID=iSee857
```

Seeing this Session ID, Tomcat dutifully retrieves the stored file, deserializes it, and executes the embedded Java code, which typically grants full remote access to the attacker.

This is about as horrible as a remote attack can get because it's dead simple to execute, requires no authentication and very little imagination. The only technical requirement is that the Tomcat server is using file-based session storage, which is common in many deployments. Also, the use of base64 encoding allows the exploit to bypass traditional security filters, making detection somewhat more challenging. And, of course, you need to know to look for it in the first place.

Wallarm detected the first attack in the early afternoon of Mar 12, Central Standard Time, originating from Poland a few days before the first public exploit was released on GitHub.
https://github.com/iSee857/CVE-2025-24813-PoC/blob/main/Tomcat_CVE-2025-24813_RCE.py

And the Wallarm folks caution about the future, writing:

> *While this exploit abuses session storage, the bigger issue is partial PUT handling in Tomcat, which allows uploading practically any file anywhere. Attackers will soon start shifting their tactics, uploading malicious JSP files, modifying configurations, and planting backdoors outside session storage. This is just the first wave.*
>
> *The reality is that reactive security—waiting for CVEs, adding Web Application Firewall rules, and hoping logs will catch threats—is a losing game. CVE-2025-24813 went from disclosure to public exploit in just 30 hours.*

That's not time for Apache's Tomcat team to get up to speed and patch, let alone test and deploy a critical update... to say nothing of having those updates deployed and servers patched.

NIST's National Vulnerability Database concurs about the severity of CVE-2025-24813 assigning it the maximum common CVSS severity rating of 9.8 and formally labeling it "CRITICAL".

The global inventory of these Apache Tomcat servers appears to be somewhere just short of around nineteen thousand installations. That's not a huge amount of global exposure, but they are likely to be running in enterprises that would qualify as prime targets.

Our takeaway here is the refrain that security is difficult and features will almost always come back to bite you in the butt.

**Severity 10/10 in AMI MegaRAC baseboard management controllers (BMCs)**

Before we leave the topic of really bad remotely exploitable vulnerabilities I should mention that the firmware security company Eclypsium discovered a remotely exploitable vulnerability in AMI MegaRAC baseboard management controllers (BMCs). The vulnerability, which is being tracked as CVE-2024-54085, received a 10/10 severity score. The reason for the maximum score is that the vulnerability allows attackers to bypass authentication and access the baseboard management controllers remote management capabilities. This would allow attackers flaw to tamper with firmware, disable security protections, and even brick devices. Over 1,000 devices with MegaRAC interfaces are currently exposed on the Internet with ASUS, ASRockRack, and HPE being the major vendors that supplied the machines.

**Google purchased Wiz Cloud Security**

We've recently covered some news involving the good work of the cloud security startup "Wiz". And due to the sound of its name I felt the need to spell it: W.I.Z., as in Wizard. In case we talk about them in the future I wanted to note for the record that they were just acquired by Google in what must have made their venture capital investors quite happy, since, as I said, this was a startup and the acquisition was the largest cybersecurity-related acquisition ever. So Google doesn't appear to be shrinking.

Google first attempted to purchase Wiz last year for a measly $23 BILLION dollars, but that deal fell through and I imagine there was plenty of disappointment to go around. But Google came back again, this time closing the deal for $32 billion in cash. The deal will need to pass regulatory review, and that might not be such smooth sailing But I have no real idea. Since I expect we'll be encountering them in the future as we have in the past, I wanted to mention that, like Mandiant, they are now a part of the Google juggernaut.

**RCS to get E2EE!**
GSMA is the GSM Association, where GSM stands for the Global System for Mobile communications. They made some news Friday before last with their announcement's headline: "RCS Encryption: A Leap Towards Secure and Interoperable Messaging" Here's what Tom Van Pelt the Technical Director of GSMA posted:

> *In my last post, 'RCS Now in iOS: a New Chapter for Mobile Messaging', I celebrated the integration of Rich Communication Services (RCS) with Apple's iOS 18, a culmination of years of collaboration across mobile operators, device manufacturers, and technology providers. Today, I am pleased to announce the next milestone: the availability of new GSMA specifications for RCS that include end-to-end encryption (E2EE) based on the Messaging Layer Security (MLS) protocol.*
>
> *Most notably, the new specifications define how to apply MLS within the context of RCS. These procedures ensure that messages and other content such as files remain confidential and secure as they travel between clients. That means that RCS will be the first large-scale messaging service to support interoperable E2EE between client implementations from different providers. Together with other unique security features such as SIM-based authentication, E2EE will provide RCS users with the highest level of privacy and security for stronger protection from scams, fraud and other security and privacy threats.*
>
> *These enhancements to support E2EE are the cornerstone of the new RCS Universal Profile release. In addition to E2EE, RCS Universal Profile 3.0 makes it easier for users to engage with businesses over RCS messaging through a richer deep link format and includes additional smaller enhancements such as improved codecs for audio messaging and easier management of subscriptions with business messaging senders. In addition, RCS continues to support a range of interoperable messaging functions between iOS and Android users, such as group messaging, the ability to share high-resolution media, and see read receipts and typing indicators.*
>
> *I would like to thank all of the contributors for their support in developing and finalising these new specifications; they represent significant progress in enabling even more of a thriving RCS ecosystem built on the foundation of secure and private messaging for the benefit of end-users worldwide.*

I took a brief look through the 90-page specification and it looks like the right people have been involved. Among other things, I noted that the work "Ratchet" appears 20 times in the document. We've discussed the use of ratchets for group messaging key distribution in the past having first encountered the term when we discussed Moxie Marlinspike's Axolotl Ratchet which he developed along with Trevor Perrin as part of the TextSecure project which was later rebranded and expanded into what we know today as the Signal protocol.

The bottom line is that it appears that the cross-platform RCS multimedia secure messaging protocol, that even Apple now supports as of iOS 18, will be obtaining strong end-to-end encryption and that it will be done correctly. I wonder what the UK and the EU will have to say about that?

**What WORLD are we living in today?**
Okay. Now, I want everyone to just listen to and contemplate this sentence which, for me at least, begs the question "What world are we living in today?" Here's the sentence that was published as a quick one-liner news blurb in a prestigious security newsletter:

> *An attacker used malicious Twitter replies to hack an AI crypto chatbot and steal over $105,000 worth of Ether.*

*"An attacker used malicious Twitter replies to hack an AI crypto chatbot and steal over $105,000 worth of Ether."* I don't even know what that means. First of all, you have to have some malicious Twitter replies, whatever those are. And those malicious replies need to be able to hack an AI crypto chatbot. What? Did those replies hurt the AI crypto chatbot's feelings? And what the hell is an AI crypto chatbot anyway? And who in their right mind would give this thing reign over a big pile of Ethereum cryptocurrency? What is wrong with people?

This podcast's listeners know that I'm more or less bullish on cryptocurrency. At least upon the fundamentals of the technology, which I have understood from the start well enough to code it up myself. But what this has all become — is utterly unrecognizable. *"An attacker used malicious Twitter replies to hack an AI crypto chatbot and steal over $105,000 worth of Ether."* Wow. Maybe I could try knitting. Is knitting still a thing?

Oh! ... and I forgot to mention that the Twitter account that perpetrated the heist, or the hack, or whatever the hell it was — the guy's Twitter account was "FungusMan". That's just perfect.

**White House seriously considering deal from Oracle to run TikTok**

The news on the TikTok US takeover front is that Oracle is the frontrunner at the moment. Politico's reporting about this contained enough interesting techie bits to make it worth sharing, particularly because there are still lots of technical questions left to be resolved and because it looks like it's what's going to be happening. Here's what Politico reported:

> *The software company Oracle is accelerating talks with the White House on a deal to run TikTok, though significant concerns remain about what role the app's Chinese founders will play in its ongoing U.S. operation, according to three people familiar with the discussions.*
>
> *Vice President JD Vance and national security adviser Mike Waltz, the two officials President Donald Trump has tasked with shepherding a deal to bring TikTok under U.S. ownership, are taking the lead in negotiations, while senators have voiced a desire to be read in on any talks, two of the people said. A third person described the White House discussions as in advanced stages.*
>
> *The people who were granted anonymity were not authorized to discuss sensitive details of ongoing negotiations publicly.*
>
> *It comes amid ongoing warnings from congressional Republicans and other China hawks that any new ownership deal — if it keeps TikTok's underlying technology in Chinese hands — could be only a surface-level fix to the security concerns that led to last year's sweeping bipartisan ban of the app. Key lawmakers, including concerned Republicans, are bringing in Oracle this week to discuss the possible deal and rising national security concerns, according to four people familiar with the meetings.*
>
> *One of the three people familiar with the discussions with Oracle said the deal would essentially require the U.S. government to depend on Oracle to oversee the data of American users and ensure the Chinese government doesn't have a backdoor to it — a promise the person warned would be impossible to keep.*

*The person told POLITICO: "If the Oracle deal moves forward, you still have this [algorithm] controlled by the Chinese. That means all you are doing is saying 'trust Oracle' to disseminate the data and guarantee there is no 'back door' to the data." If the algorithm isn't entirely rebuilt by its U.S. owner, or if TikTok's Beijing-based parent firm ByteDance retains a role in its operations, it could retain vulnerabilities that could be exploited by the Chinese government.*

*The data security company HaystackID, which serves as independent security inspectors for TikTok U.S., said in February that it has found no indications of internal or external malicious activity — nor has it identified any protected U.S. user data that has been shared with China.*

*Spokespeople for Oracle, TikTok, ByteDance and the White House did not respond to requests for comment. The deal is being billed as a "Project Texas 2.0," in a nod to a previous agreement between TikTok and Oracle to relocate American users' data to servers in Texas and block ByteDance employees in China from accessing it, according to the first person. But that agreement, which also required Oracle to review TikTok's source code to determine its safety, failed to assuage congressional and Biden administration concerns that the app is being used by China as a spying and propaganda tool.*

*The tech-focused outlet* ***The Information*** *reported Thursday that Oracle is a "leading contender" to run TikTok, with ByteDance preferring it for the role. The details about the White House's approach and the seriousness with which White House officials are considering the proposal have not previously been reported.*

*It comes as Trump stares down an April 5 deadline to secure a new owner for the Chinese video-sharing company after he signed an executive order in January delaying enforcement of Congress' ban on the app for 75 days. The app briefly went dark for about 12 hours in January after TikTok's parent company ByteDance failed to meet the deadline to sell its stake and the Supreme Court upheld Congress' ban.*

*Vance, during an interview with NBC News on Friday, said he was hopeful a TikTok deal would be reached by the early April deadline. Last week, Trump said that his administration was in talks with "four different groups" about a deal.*

*Trump told reporters in January that he was open to Oracle founder and executive chairman Larry Ellison buying TikTok. Ellison is a longtime Trump supporter, and he's part of so-called Project Stargate, a $500 billion AI infrastructure initiative that also includes OpenAI, SoftBank and MGX.*

*While Trump during his first administration sought to ban TikTok over national security concerns, he embraced the app last year on the campaign trail. In December, he told throngs of young conservative supporters at a Turning Point rally in Phoenix that he has "a warm spot in my heart for TikTok" because of the outpouring of support he received from younger voters in the 2024 election.*

*It's unclear whether the deal the White House eventually reaches will satisfy China hawks on the Hill, though they may have little power to complain. Trump's executive order extending the initial deadline — in the face of concerns from GOP lawmakers and legal experts about the order's legality — showed his willingness to defy Congress' will. And the decision on whether ByteDance sells TikTok or licenses its use by a U.S. company ultimately rests with the Chinese government.*

*Beijing wants to protect TikTok's monopoly access to its user data and is hostile to any suggestion that Chinese firms bend to the will of suspicious foreign governments. Over the past year, authorities in Beijing and in the Chinese embassy in Washington have mostly dodged questions about the status of possible talks for the purchase of TikTok by a non-Chinese firm.*

*What little Beijing has said about that possibility hasn't offered much hope that it's in favor of such an agreement. The Chinese government "will firmly oppose" any forced sale of the company and require ByteDance "to seek governmental approval in accordance with Chinese regulations" for any potential foreign ownership deal, a Chinese Commerce Ministry spokes-person told reporters in March. That same month a Chinese Foreign Ministry spokesperson accused Congress of "resorting to hegemonic moves" to try to take control of the app. In January, the Chinese government deployed more conciliatory language about a possible TikTok sale but offered no clues on whether it would approve such a deal.*

*Any such transactions "should be independently decided by companies in accordance with market principles," a Chinese Foreign Ministry spokesperson said in January.*

**23andMe going under**

Two days ago, on Sunday March 23rd, the original personal genomics company 23andMe filed for protection under chapter 11 of the bankruptcy act. Their press release had the headline: *"23andMe Initiates Voluntary Chapter 11 Process to Maximize Stakeholder Value Through Court-Supervised Sale Process"*. I'm mentioning this here because, from a personal privacy standpoint, now might be a good time for anyone worried about the future of any of their genetic data being held by 23andMe to delete it from their databases and to close their account. As a founding member of 23andMe, I just did exactly that:

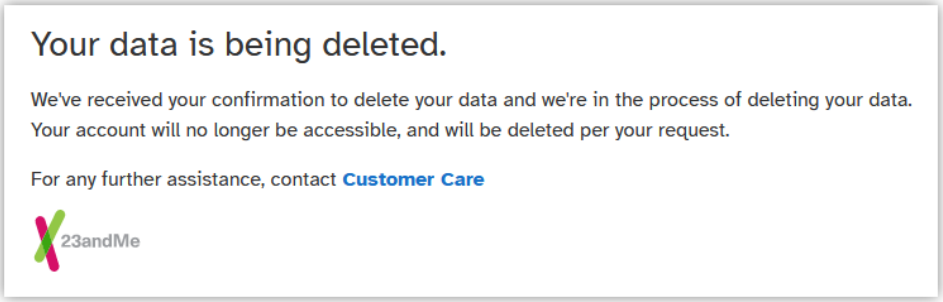


Since it took me some poking around their website, I recorded the process to make it easier for any who might wish to do the same. I spit in their test tube long ago, and I'm not in a panic about it. But given that they're going under and someone I don't know will be purchasing their assets for pennies on the dollar, leaving my genetic data behind in their database seems unlikely to do me any good at this point.

So I logged in, selected "Settings" under my shadow head and shoulders in the upper right. Once that page came up (which took awhile, so I may not have been alone), scroll to the very bottom of the page to the "23andMe Data" section. Then click the "View" button. I noted that the View page has a clean looking URL that would also take you directly to the page you need. It's: https://you.23andme.com/user/edit/records/ Or, after logging in, you could use the GRC shortcut link I created to jump directly to the sayonara page: https://grc.sc/byebye.

**The White House says not to fire any Cyber Guys**
And, finally, in some good news for cyber security professionals, the White House administration has reportedly told federal agencies to avoid firing any cyber guys. Here's part of what Reuters wrote under their headline: *"White House instructs agencies to avoid firing cybersecurity staff, email says"*:

> *According to an email seen by Reuters, the White House is urging federal agencies to refrain from laying off their cybersecurity teams, as they scramble to comply with a Thursday deadline to submit mass layoff plans to slash their budgets. Greg Barbaccia, the United States federal chief information officer, sent the message Wednesday, in response to questions about whether cybersecurity employees' work is national security-related, and therefore exempt from layoffs. He wrote in the email to information technology employees across the federal government which has not been previously reported: "We believe cybersecurity is national security and we encourage Department-level Chief Information Officers to consider this when reviewing their organizations."*
>
> *Describing "skilled cyber security professionals" as playing "a vital role in mission delivery and information assurance", he said, "We are confident federal agencies will be able to identify efficiencies across their non-cyber mission areas without negatively affecting their agency's cyber posture."*

As part of the downsizing that Trump and Musk have controversially been engaged in recently, CISA had more than 130 positions cut. We've been talking about CISA more and more often for the past few years since they've objectively been doing a really terrific job, which is astonishingly unusual for anything within the government bureaucracy. I certainly never expected CISA to amount to what it has. So I've been hoping that CISA would survive and remain as highly functional as they have been. And to that end, there was some recent news that those jobs were being reinstated. So that's reassuring. We need CISA. They've really been implementing some terrific policies and creating needed requirements for the cybersecurity of federal agencies.

**AI project failure rates are on the rise**
An interesting piece in Cybersecurity Dive caught my eye. It was a report that said that AI project failure rates were on the rise. I thought that was interesting. It suggests that just slapping a ***"Now even more better with AI!"*** label on anything and everything may not always produce a win. My guess, though, about the reason for failure **rates** rising is mostly the explosion in all of those labels being hastily added.

Still, it was interesting that according to a report from S&P Global Market Intelligence, based upon a survey of more than 1,000 responding enterprises in North America and Europe, the share of businesses scrapping most of their AI initiatives increased to 42% this year, up from 17% last year. Again, I'm sure largely because so many more were trying.

The average organization scrapped 46% of AI proof-of-concepts before they even reached production. The surveyed enterprises cited cost, data privacy and security risks as the top obstacles. (I wonder whether they heard the news about that AI crypto chatbot?) Anyway, at this point AI adoption is predominantly found in IT operations, followed by customer experience workflows and marketing processes.

So it appears that the initial "AI Everywhere" euphoria is quickly coming back down to earth and closer to reality. I'm sure not letting it get anywhere near SpinRite! ... speaking of which ...

# Listener Feedback

**Ken**

*Hi Steve: Ken here (65 years old), Canadian trucker for 40 years. I just want to say thank you for your dedication and enthusiasm in the tech world and the beautiful things you have contributed to tech. I just bought SpinRite recently and it's a total game changer. I ran it on my current machine and it tuned up my SSD's like crazy..... amazing software, thank you. I build computers and repair them and recently a buddy of mine dropped off an old Windows 7 machine that was in a closet for 7 years. He wanted the old pictures from it, of course, I managed to get it to boot and got all his old pics and transferred them to a new rig I had ready to go. I ran SpinRite of course ...and now that old beast runs like a champ.*

Thank you for your report, Ken. The best thing about SpinRite is that, aside from it being the miracle that has largely provided for my life, I get to hear about how much its use helps people, and nothing beats that.

**Tom**

*Hi Steve, Now that ublock origin is no longer supported in chrome, I'm going to start using firefox. I've exported my bookmarks from chrome to firefox, but I'll likely be using both browsers, at least for the time being. Do you know of any browser extension that mirrors favorites between chrome and firefox? If I make a change to any bookmarks while I'm using chrome, I'd like for those changes to sync to my chrome bookmarks. Thanks, Tom*

That's a terrific question. I suppose I've become so accustomed to only using a single browser platform at a time, and just assumed that each would have its own native and closed ecosystem, that I never considered wanting or needing cross-platform synchronization.

So, spurred by Tom's question, I poked around and found a very nice looking 3rd-party cross-platform extension for Chrome and Firefox desktops as well as Android, called "xBrowserSync" https://www.xbrowsersync.org/ And, boy, these guys sure are saying all the right things. A snippet from their site says:

*xBrowserSync is a free and open-source alternative to browser syncing tools offered by companies like Google, Firefox, Opera and others. The project was born out of a concern for the over-reliance on services provided by big tech, who collect as much personal data as they can and have demonstrated that they do not respect their user's privacy. Now, with the proliferation of open-source code and projects it's easier than ever to create tools and services that allow users to take back control of their data!*

*xBrowserSync respects your privacy and gives you complete anonymity. No sign up is required and no personal data is ever collected. To start syncing simply download xBrowserSync for your desktop browser or mobile platform, enter an encryption password and click Create New Sync! You'll receive an anonymous sync ID which identifies your data and can be used to access your data on other browsers and devices.*

*xBrowserSync does not only sync but also enhances your productivity by enriching your native browser bookmarks with the addition of descriptions and tags, and an intuitive search interface enables you to find, modify and share bookmarks quickly and easily. xBrowserSync even adds descriptions and tags to new bookmarks for you automatically. And don't ever worry about*

*losing your data thanks to the included back up and restore functionality.*

*The xBrowserSync desktop browser web extension syncs your browser data between desktop browsers. It works with the browser's native bookmarking features so you can keep using the native tools whilst always staying in sync. If you like to organise your bookmarks into folders don't worry, xBrowserSync respects your bookmark hierarchy and syncs it across your browsers.*

That sure sounds like exactly what Tom is looking for, and it's from folks who clearly share the spirit and philosophy we'd like them to have. After reading Tom's note and running across that xBrowserSync extension, I sent this all back to Tom. Not long after he replied:

*Thanks Steve. I will look into this a bit more, but when I clicked to download for chrome, I'm taken to the chrome web store which shows this: "This extension is no longer available because it doesn't follow best practices for Chrome extensions." Thanks, Tom*

Wow. That sure sounds like the Chrome folks don't like the whole idea of cross-platform browser synchronization which would have facilitated Tom's gradual migration away from Chrome. Given the way the xBrowserSync people described their intent and philosophy, it would certainly appear to be 100% the sort of practices that anyone would feel good about endorsing. It would be disappointing if this was purely an anti-competitive move by Google.

While the fate of any other similar extensions might be similar, for what it's worth while searching around I did find several others. But again, if Google's Chrome is threatened by these, then any of them might be a lost cause.

**BackGhost**

*I found your comments on the state of vendor support for old and outdated hardware intriguing and wanted to add more insight into what is a very complex issue. As I work for a service provider that is also a manufacturer of networking gear and often see both sides of this issue.*

*Hardware manufacturers deal with the same software and hardware EOL/EOS (End of Life and End of Support) issues as customers, just at a micro level. Every ASIC/CPU/IC has a lifetime and its own software with a lifetime. When vendors have to support more products from a software and hardware standpoint, it costs the vendor more. The vendor can and often does charge more for this support of old gear, but at some point, the cost of support will outweigh the cost that can be charged to a shrinking set of customers. Vendors will often discount or offer trade-ins for old gear to encourage customers to upgrade to new gear. Luckily, the vendors (well, the big iron guys) will give advance notices of EOL/EOS, and the sales team is always eager to engage the customer on new sales opportunities. As service providers we struggle with the never ending notices of EOL/EOS of gear and will often have to fight for capital to do upgrades and or replacements. These efforts will be taken on based on business objectives and risk, etc. and leads to the never ending dance between the CTO, CFO, Sales, Product development.*

***The service provider side:***

*Hardware manufacturers will always EOL (End Of Life) equipment and often give notice well in*

*advance. Larger companies that sell "big iron" will give notice years out. For example, Juniper (off the top of my head) provides 3 years for hardware support and 1+2 years on software support after the hardware is no longer supported for replacement support. So, there is normally "plenty" of time for planning for obsolescence and replacement. Of course, these replacement plans are driven by business goals, which leads to point 2.*

*The CIO/CFO battles are the norm, and this battle is complex at best. Do we update now, later, or never? Do we roll the dice? Are we doing a new build somewhere else that has our focus? These are endless. Just to say, it is complex.*

*The other side of this equation is the hardware manufacturer side, and this is what drove me to send this feedback:*

***Hardware Vendor side:***

*On the Hardware support side: (a) Discrete components (IC, chips, etc.) can no longer be sourced. (b) Discrete component replacement causes board redesign, and the cost of redesign is too high. (c) Discrete component software support is EOL due to the manufacturer EOL of the IC. The IC library is no longer supported due to EOL on the software support. (d) The new replacement product is just cheaper, better, and faster. Why keep the old one around given its install base? (This too is complex and often political. You don't want to upset a long-time big customer with a hardware upgrade.)*

*And on the Software support side (just a few): (a) See hardware support 1c, as this is part of the software chain. (b) OS and supporting software are no longer supported by vendors. (c) Newer upgraded replacement hardware uses different software for various reasons and thus is not compatible with old hardware. This causes a complete new software support, development, and test chain. (d) The cost of support is higher than the customer can sustain and can drive the customer to find other solutions. Like the hardware side, this is complex and often political. (e) Software licensing has a lifetime, limited in volume, developer seats, etc., that forces an EOL action.*

***Planned obsolescence:***

*(a) Most manufacturers can predict when a given bit of hardware they sell will be EOL/EOS and often balance the above two major points to give a date of EOL/EOS. (b) Hardware roadmaps will take this into account and offer new, better, and often cheaper products to the customer and better cost management for the manufacturer.*

*These are just a few items, and I have many stories from the field on the service provider and manufacturer side and the confluence between the two. These are some good stories that are often shared between peers over cappuccinos, a good mead, or beer.*

*I find the podcast intriguing and thought-provoking. Obviously using my non-work related email to send this comment!  All the best*

I thought this person's comments were worth sharing. For one thing, I would never expect ongoing ***hardware*** support for any device beyond the manufacturer's original commitment. It it might be available, fine. Things like power supplies can be somewhat generic and might be easily replaceable. But if, for example, ports die on an expensive router or switch that's out of warranty then the calculus is entirely different and the conversation with the CFO is very different: "The mission critical device just died, we're currently limping along and we need it

replaced ASAP!" That's not the conversation I hypothesized last week. I do really understand that maintaining old software has a decidedly non-zero cost. But the point I was making last week was that revenue is being left on the table. The manufacturer hopes that a lack of ongoing support will force their customers to move to newer equipment. But the reality is, most will remain with out-of-warranty and out-of-support equipment and suffer the potential consequences.

**Dan Linder**

*Hi Steve,*

*In security now, episode 1017 you made a comment about a Juniper router being unsupported and vulnerable, and then a hypothetical conversation between a CIO and a CFO about replacing that otherwise hardware just because it was out of support.*

*I too have some experience with US department of defense rules, and one thing that I haven't heard you discuss on the show are the stig documents, STIG stands for "Security Technical Implementation Guide".*

*The stig document is a series of checks (or control) and actions to take on a specific system that can harden it to some degree to mitigate threats to its overall security. Each control is given a category 1, 2, or 3 rating, with "Cat-1" being the most important controls to implement. Within each control there are some check text steps and corresponding fix text steps which list a simple command or action to take to validate that the control is in place, and if not, what can be done to enable it.*

*While the STIGs give a specific fix text to implement, most security organizations that review the application of these STIG controls allow for additional/external controls that will mitigate a specific problem if it can't be addressed with the fix text suggested. For instance, if an insecure system is being used, but it is only used in an air gapped environment, only accessible by a small number of people already vetted and trusted, they might be willing to overlook a Cat-1 finding.*

*In all the STIGs I have worked with, they all have a security question which requires confirmation that the system being secured can still receive updates from the manufacturer. If the company in your example was applying and enforcing the STIGS as written, then the CIO has quite a bit of leverage to go back to the CFO to get this system replaced.*

*I hope you can find time in a future episode to give a brief talk about the stig documentation, and some of the potential for securing anyone's environment regardless of government affiliation.*

*Thanks, Dan Linder*

That's great information, Dan. Giving CIO's all the ammunition they can get to justify the non-zero cost of ongoing SUPPORTED maintenance of a constantly aging infrastructure is what we want. Now we need similar strong policies from CISA. I'm unsure how they would be forced. But it would help with the necessary change in culture for the CFO to appreciate the real risks of continuing to use past end-of-life technologies.

(Perhaps mentioning that the current equipment is not quantum safe might work?)
And on THAT note! ...

# The Quantum Threat

We love showing up for this podcast every week, which, after all, we've been doing for nearly 20 years. And, much as I would dearly love to be, I doubt we'll still be here the day a quantum computer first cracks actual, working-strength, public key encryption – but, **BOY!**, ***that*** would be a terrific ripping podcast! Nevertheless, that day still seems quite a ways away. What is clear, regardless of exactly **when** that may happen, is that the future threat to encryption is very real and is becoming more real every day. Through the years of this podcast, we've all become students of the history of computer security. And one lesson we've all learned together is just how very very long it's going to take to wash all of the old pre-quantum crypto out of our existing systems. That all leads to the simple and incontrovertible conclusion that there's no time like the present to begin.

Last Tuesday, Hewlett-Packard's "Threat Research" group posted a terrific piece titled *"From False Alarms to Real Threats: Protecting Cryptography Against Quantum"*. That's what I want to share today. In their opening, they make some great points that are well worth appreciating:

> Quantum computers could break asymmetric cryptography, which would be catastrophic for society's digital infrastructure. Quantum computers powerful enough to break cryptography do not exist today, but the threat of one being created steadily advanced in 2024. With multiple quantum computing technologies overcoming development obstacles, the security community is now more sure than ever that sufficiently powerful quantum computers will come. Some think it could be ten years, but with the speed of recent innovation, an unexpected breakthrough could accelerate that. This has created a significant security risk because we rely on protections for a long time and need them in place before threats arise.
>
> Since we last wrote on this topic a year ago, authorities around the world have increased efforts to urge organisations to start migrating systems to quantum-resistant cryptography. Critical industries are especially advised to mitigate these quantum risks given they are high profile targets. Particular priorities for migration include sensitive data vulnerable to capture-and-decrypt attacks, and protections rooted in hardware. Without upgraded protections at the hardware and firmware foundation, quantum attackers can compromise devices even if the software running on the hardware is quantum-resistant.
>
> 2024 also saw several false alarms of quantum breaks to cryptography. We expect that to become a trend as innovation in quantum computing progresses. What we have seen is that such false alarms will elicit panic from some, but only complacency from others. But they also proved useful in raising the conversation about readiness and an understanding of the consequences of a real alarm. In short, we must stay vigilant and prepare for the real threat.
>
> Over the last year, we at HP also made progress to protect customers from the threat of cryptography being broken by quantum computers. Last year we announced the world's first business PCs to protect firmware integrity against quantum computer attacks. Today, we are announcing the world's first printers to protect firmware integrity against quantum computer attacks. These security innovations demonstrate our dedication to safeguarding our customers against future threats.

They then quoted Boris Balacheff, the head of the HP Security Lab, an HP Fellow and Chief Technologist for Security Research and Innovation. Boris said:

*"As innovation progresses towards more powerful quantum computers, it is urgent to prepare for the threat this represents to the asymmetric cryptography we depend on in our daily digital lives. This starts with migrating systems that cannot be updated easily once deployed. After the introduction of quantum-resistant firmware integrity protection in PCs last year, today we are announcing the launch of printers with similar capability to protect against future quantum computing threats. We continue with our commitment to lead the way with endpoint security innovation, and keep our customers safe into the future."*

This is not something we've focused upon or talked about previously. And of course they're correct. As we know, all of the secure booting technology we have today is based upon the motherboard's firmware verifying the digital signatures of the software that the motherboard's UEFI firmware first loads. And **all** of that secure boot technology is currently pre-quantum. It's embedded into the hardware with technologies such as the TPM, Trusted Platform Module, that dates from 2003.

Listening to what HP has to say here really serves to put a much finer point on this looming issue. I've edited the piece to remove HP's non-technical self-promotion and for length; but there's a great deal of great information here. They wrote:

*In the past 12 months, the cryptography and security community has experienced heightening concern over the progress of quantum computing. The last year has been marked by key developments in quantum computing technology, as well as multiple instances of false alarms over potential quantum breakthroughs that put cryptography at risk. Although these alarms were ultimately disproven, when considered alongside genuine advancements in quantum computing, they highlighted the fragility of society's digital infrastructure. A sufficiently powerful quantum computer could break much of the cryptography relied upon globally. Given how fundamental cryptography is to security everywhere, a quantum computing breakthrough before the world is ready would jeopardise security. It could allow attackers to run riot across our digital infrastructure – giving them freedom to access network services, takeover devices, steal blockchain assets, decrypt sensitive data, and more.*

*In reaction to these advancements, there has been an increased sense of urgency to fortify cryptography, driven by technical authorities and experts. This urgency has led to accelerated timelines and new policies to address the looming quantum threat. Against this backdrop, the security community has intensified its preparations. Academia, standards bodies, governments, and industry are collaborating and making concerted efforts to migrate technologies to being quantum-resistant.*

*In this blog post, we discuss two false alarms that percolated through the community over the last year and what we learned from them. We explore the current state of the quantum computing threat to cryptography and how the community is preparing a response.*

*The first alarm took place in April of 2024 during the NIST 5th PQC (Post Quantum Computing) standardization conference, which had convened to discuss cryptography designed to withstand quantum computer attacks. The trigger for the alarm was an academic paper – newly published and not yet reviewed or corroborated – describing a new quantum computer attack that could have been effective at* ***breaking*** *the new post-quantum cryptography the technical community had been working on for almost a decade. This cryptography was meant to become a global standard to protect digital infrastructure, should quantum computers break traditional asymmetric cryptography like RSA and most Elliptic Curve Cryptography (ECC).*

*A claim it was broken was shocking and would leave the quantum-resistant migration in disarray, if confirmed true. Speculation about the paper, entitled "Quantum Algorithm for Solving Lattice-Based Cryptosystems", lit up our technical social media networks. One of our team was at the conference. While the talks continued and the audience listened attentively, attendees gradually started to form small huddles, trying to make sense of the publication. Remarkably, no one was sure the paper was incorrect. Most hoped it probably was incorrect, but at face value it was convincing – presenting a credible nine-step algorithm that put quantum-resistant lattice-based cryptography in a very precarious position.*

*For eight days, there was furious analysis among cryptographers and quantum computation experts. But with very few people who can claim to be experts in both fields, many researchers wrestled with analysis beyond their areas of expertise. A Discord community sprang up, crowd-sourcing a comprehensive analysis and triage of the paper's claims. This intense assessment-phase ended when two researchers found an inconsistency in the final step of the algorithm. The paper's author engaged with this critique and confirmed the final step had an irreconcilable error.*

*And thus, the community breathed again. But for an entire week, the community responsible for developing the cryptography that will protect much of our digital lives into the future had seriously considered the possibility that they could have got it wrong. Because this was so technical and didn't impact the cryptography we use currently, the news did not make the broader security community panic – and the doubt didn't last long enough within the cryptography technical community to gain momentum and spread.*

Our podcast listeners may recall that we did touch the fact of this having happened at the time. HP continues:

*The second moment of 2024 when the broader security community thought that cryptography was broken was also triggered by an academic paper. The paper, "Quantum Annealing Public Key Cryptographic Attack Algorithm Based on D-Wave Advantage", was published in May 2024 in the Chinese Journal of Computing. This false alarm caused more widespread uncertainty and panic in the technical community and beyond, with several reports stating incorrectly that some researchers were able to break RSA encryption using a D-Wave Advantage quantum computer.*

And, again, that news made it into this podcast – because it would be difficult to overstate just what havoc would ensure if this were to be true.

*With a general audience unable to assess the original paper (only the abstract was published in English), the reports generated real anxiety. However, there was little credibility in the claim that RSA had been broken, and expert consensus rapidly emerged. With a bit of scrutiny, it was established that the researchers had only broken a very small scale, simplified RSA, and their solution did not scale to the kind of numbers used for security and was therefore not a credible threat.*

*Once again, after a week or so, concerns about pre-quantum cryptography having been broken were largely quelled. However, for several months afterwards, incorrect reports still appeared, sparking fresh waves of concern among those who had missed the initial reporting.*

*One benefit of these events is that they test the security community's preparedness for the sudden removal of some fundamental underlying cryptographic primitive. From that perspective, these alarms have been like the safety briefing before an airplane flight – forcing the community to grapple with what to do in the worst-case scenario. If the event were real, are we ready? What preparations should be in place, and are they?*

*The fact that a broad audience was alarmed tells us that there is a growing understanding of the critical impact of the quantum threat, and that action will increasingly be called for. The successful resolution of these incidents underscores the importance of a measured and collaborative approach to evaluating cryptographic research, for the community has shown it can be relied upon to robustly evaluate these complicated ideas. Unfortunately, analysing such academic papers is inherently complex, requiring expertise that is rarefied and spans multiple fields – cryptography, mathematics, quantum algorithms, quantum computer engineering and physics. So, we should anticipate regular moments of doubt in the security of our cryptography and have the patience to wait for assessment before panic-induced reactions.*

*One day, there could be surprise news, or even a significant rumour, of a real breakthrough. Rather than panic, we should instead ensure we are prepared and have put in place quantum-resistant protections – starting with our priorities.*

*This said, there is also concern that too many false alarms related to quantum computing breakthroughs could eventually lead to a sense of complacency and inaction. This might cause people to believe the quantum threat is not yet a serious concern. If too many incidents lead to unwarranted panic, a genuine threat might be ignored as just another false alarm when it finally arises.*

What becomes clear is that where we need to be, and as soon as is practical, is at a point where we're no longer reliant upon classical pre-quantum crypto so that the eventual announcement of a true breakthrough is met with a yawn and a shrug. So, where exactly are we today? What is the current true level of alarm we should be feeling? HP addresses that, writing:

*With so many possible quantum breakthroughs to be assessed, and uncertainty about what is credible, it can be difficult to understand the landscape of quantum computing and separate fact from fiction. Let's take a closer look at the reality.*

*To gauge the true alarm level, we should examine the progress in quantum computing technology. Over the past year, there has been impressive advancement in several technologies, with multiple promising pathways emerging. Even if some fail, others may succeed. Compared to a year ago, large-scale quantum computing now seems more likely. We look to experts to quantify this likelihood.*

*The Global Risk Institute's 2024 report highlights a* ***"significant chance"*** *of a quantum threat emerging by* ***2034****, posing an* ***"intolerable risk from a cybersecurity perspective"****. Nearly one third of the 32 experts surveyed estimate a 50% or greater chance of quantum computers breaking cryptography by 2034, with an average estimate of 27% – the highest in the six annual surveys conducted so far.*

Let me repeat that:

*Nearly one third of the 32 experts surveyed estimate a 50% or greater chance of quantum*

*computers breaking cryptography by 2034. The average estimate of a quantum computing break by 2034 was 27%.*

*To summarise recent the change, the report states: "The progress in the last year has induced many people both within and outside the quantum research community to realize that the quantum threat may be closer than they thought." The German Information Security authority, BSI, recently updated their comprehensive assessment of quantum computer technologies. The report concludes that, due to major roadblocks being* ***resolved****, quantum computers are likely to break cryptography within at most 16 years but recognises that new developments could lead to a breakthrough as soon as a decade.*

*Progress has been made not only in various quantum computing candidate technologies, but also in aspects like stability, scale, inter-connectivity, and operating software.*

*Stability is a major challenge for current quantum technologies, as they do not hold their state for long before deteriorating. Reducing noise and using effective error-correction – where more errors are corrected than introduced – is crucial for long-term stability. Demonstrating this effectiveness is a milestone that has been achieved by four technologies: Superconducting Transmons, Ion Traps, Neutral Atoms, and Color Centers.*

*Sizes of systems have increased as production processes mature, with Google announcing their 105-qubit Willow, IBM introducing the 156-qubit Heron along with a roadmap for processor scaling, and Microsoft and Quantinuum upgrading the H2 Trapped Ion processor to 56 qubits.*

*The stability and size of the relatively new Neutral Atom technology, whose key elements were only demonstrated as recently as 2022, has also shown a massive improvement with potential for acceleration. The QuEra start-up that came out of this research has just this February been backed with a $230M investment, giving an indication of the high interest in this approach. Of very recent note, a new technology with greater natural stability – the topological qubit – has been demonstrated for the first time as a proof of concept by Microsoft who claim the technology offers a "clear path to fit a million qubits on a single chip", which would be needed for scaling.*

*Advances in inter-connecting quantum states between different chips are starting to show promise for enabling the distributed quantum computation needed for large quantum computers. Additionally, an ecosystem of organisations are developing the necessary developer tools and software stack for operating quantum computers and creating quantum programs. This stack, like the classical computation stack, ranges from physical machine instructions to higher-level programming languages, allowing specialists to effectively use their expertise and enhance progress.*

*Given all these advancements, Scott Aaronson, a quantum computing expert, recently said he believes that "the race to build a scalable fault-tolerant quantum computer is actually underway". His position on the urgency of addressing the quantum threat to cryptography has shifted from "maybe" to* ***"unequivocally, worry about this now. Have a plan."***

*In summary, in just the past year, breakthroughs in quantum computing have strengthened the consensus that quantum computers capable of breaking today's cryptography may become feasible soon. It may only take a surprise acceleration from one of the promising technologies to break cryptography in less than a decade. Therefore, it's crucial to assess our preparedness and take action to ensure we are fully ready.*

And then HP notes almost needlessly: *"Migrating Quantum-Vulnerable Cryptography is on a Whole New Level Compared to Patching a 0-Day Vulnerability"* Although I'm sure listeners are aware that we're talking about a sea change that requires us to scrap everything we've built, it's worth hearing HP out on this. They write:

> *It is tempting to think the problem of fixing quantum-vulnerable cryptography is like patching a zero-day vulnerability in code. However, this analogy under-represents the scope of the quantum threat. A zero-day vulnerability is an error in a specific sequence of computer instructions in a specific program or library, which can typically be identified and then patched. Even if the error occurs in a pervasively common library, such as the Log4j vulnerability, it is still fixable by deploying a patch.*
>
> *Unlike a 0-day, the quantum threat does not apply to a specific sequence of computer instructions but instead applies to* ***all implementations of vulnerable asymmetric cryptography.*** *These implementations vary widely, potentially manifesting in millions of different code sequences. When quantum computers become viable, each of these will need replacement individually, by upgrading the cryptographic algorithms and keys used, requiring a global effort and collaboration by security practitioners, business leaders, and cryptographic experts.*

You know, the more I think about it, the more I think I'm glad that this podcast will probably not be around to see this disaster befall humanity. Really! Given the reluctance to change that we've witnessed throughout the past 20 years, what chance is there that we're going to be least bit prepared for this? We're talking about replacing everything – and doing it even while it's not obviously necessary that it needs to be done at all. And remember that security is only as strong as the weakest link. Who's not going to have some old webcam, lightswitch, thermostat or router lying around that continues relying upon pre-quantum crypto? HP wrote:

> *This process of patching has already started and is part of the migration to quantum-resistant cryptography that the security community is currently undertaking. But how should organisations be responding?*
>
> *Across government, industry, academia and standards bodies, mechanisms to protect against quantum attacks are being put into place with some urgency. Our advice is to start by inventorying what would be vulnerable to quantum attackers. Then prioritise what needs migrating and protecting first. The most urgent priorities for most organisations include:*
>
> - *Protecting data with long-term confidentiality requirements*
> - *Protecting long-lived systems by upgrading cryptography in hardware*
>
> *The cost of upgrading hardware is expected to be significant. In July 2024 the US Office of the National Cyber Director published a report estimating the total cost of quantum-resistant cryptography migration for prioritised US government systems between 2025 and 2035 at approximately $7.1 billion. In their calculation, they specifically call out that migrating the cryptography hardwired into hardware or firmware would constitute a significant portion of that overall cost.*
>
> *Government authorities are uniquely positioned with expert insights and the responsibility to protect national assets. Understanding their strategy and policies for critical systems and infrastructure should help any organisation plan for migration with appropriate urgency.*

And let's hope we have a vital and functioning CISA to keep this on the forefront of everyone's mind. HP continues...

> *Let's start with the US, who have a comprehensive plan and set of actions in place. In 2022, US authorities established a tempo for migration. This has led to all federal agencies planning, taking inventories, and reporting on progress annually. A timetable to migrate National Security Systems was also established,* ***with all new acquisitions from 2027 needing to be quantum-resistant****, and all non-migrated products to have been phased out by the end of 2030.*

Wow. That's great!

> *Migration of firmware signing is prioritised as even more urgent, with migration of firmware roots of trust – the firmware integrity protections in hardware – expected to be "implemented for some long-lived signatures in 2025". Since 2022, authorities have put in place guidance, including a guide published by CISA, NSA and NIST, and organised outreach to help engage and ready industry. Most recently, the Executive Order on "Strengthening and Promoting Innovation in the Nation's Cybersecurity" of 16 January 2025, further emphasised the urgency to migrate. It specified that when procuring products, federal agencies must require quantum-resistant cryptography when it is widely available in a product category and require quantum-resistant protection in networks "as soon as practical".*
>
> *Alongside this, NIST recently released its draft plan to deprecate classical asymmetric cryptography – RSA and relevant ECC – from the end of 2030 and entirely disallow it for security purposes after 2035. Assuming this plan is confirmed, this will be highly influential in establishing migration urgency, because it means there is an end date within the lifetime of many current systems. Even during 2031-2035, data owners will only be able to use quantum-vulnerable cryptography by exception, where they evaluate and accept the risk.*
>
> *Beyond the US, the Australian Cyber Security Centre (ACSC) is also setting an urgent timeline for migration. The ACSC recently updated its Cryptography Guidelines for government and industry to disallow quantum-vulnerable cryptography after 2030.*
>
> *In Europe, the security authorities of the UK, France, Germany, the Netherlands, Sweden, Norway, and Switzerland, all urge preparation and are giving increasingly comprehensive guidance on how to migrate and prioritise. In April 2024, the European Commission recommended establishing a strategy to migrate public services and critical infrastructures as soon as possible. Building on this, in November 2024, 18 EU Member States issued a Joint Statement urging nations to make the transition to quantum-resistant cryptography a "top priority" and protect the most sensitive data "as soon as possible, latest by the end of 2030."*
>
> *The last 12 months have seen an intensification of the calls to migrate by national authorities. This underlines the need to act: assess cryptography dependencies, plan and prioritise for migration, and start to migrate priority assets.*
>
> *The heightening of the quantum threat to cryptography and the intensification of national calls to action during the last year have fortunately been met with significant progress in the range and availability of mitigation solutions. New quantum-resistant cryptographic algorithms were released as NIST Standards last year to celebration of government, academia and industry – following a collaborative selection process spanning nearly a decade. These new algorithms offer quantum resistance suitable for general use in protocols and applications. They also*

*complement existing standardised quantum-resistant hash-based signatures suitable for special purposes, such as code signing. With this suite of standards, it has now become possible for industry to migrate in many scenarios.*

*Standards capture community consensus and security best practice, while enabling interoperability between different elements across a system. As such, standards are a crucial part of industry migration to quantum resistance. From standards that define new cryptographic algorithms, through to protocols that use these algorithms and applications that adopt them, the community is carefully and steadily integrating quantum resistance into the technology stack and making resistance available to customers in products.*

*This is why collaborating with other vendors and participating in standardisation efforts is essential. Notably, HP is engaged in NIST's National Cybersecurity Center of Excellence (NCCoE) Migration to Post-Quantum Cryptography project. This NCCoE project was convened to bring industry and end-user organisations together to help solve the practicalities of quantum resistance adoption and transition.*

*To stay ahead of the quantum threat to cryptography, we cannot afford to take a "wait and see" approach. At HP, our strategy is to prioritise quantum resistance from the hardware up and securely migrate from there. When prioritising and planning what protections to migrate, it is crucial to consider the cost, effort and difficulty of engineering the change. Migrating hardware – and the solutions baked into hardware – often requires changes to physically-engineered parts, which can be slow and needs a lot of forward planning, sometimes years ahead.*

And even though it's a bit of a sale pitch, it's worth looking at what HP is doing as a case study:

*Last year, we introduced PCs with quantum-resistant protection of firmware integrity designed into hardware. Our quantum-resistant hardware foundation defends against a quantum computer attacker forging the signature on malicious firmware and taking over the device. Today, we announced the launch of the world's first printers designed to protect firmware integrity against quantum computer attacks. These new printers safeguard the integrity and authenticity of low-level firmware against quantum attackers, with firmware protection integrated into the hardware.*

*Without this quantum-resistant foundation, a quantum attacker could run their own code at the most privileged level and totally compromise the security of the printer – handing an attacker control of the device and access to all its data. The high impact of a successful attack and the long lifespan of modern printers were key to why we have prioritised making our printer security foundation quantum-resistant. This upgrade also provides an anchor for further quantum-resistant updates to printer software.*

That makes a lot of sense since we've seen how printers can become the home to advanced persistent threats. They are always on and are well connected to the enterprise's network. And HP is certainly right about them being set and forget ... and then being taken for granted.

They conclude, writing:

*The threat quantum computers pose to cryptography has steadily advanced this past year, creating an unacceptable security risk to the algorithms fundamental to securing our digital*

*lives. It would be devastating if these cryptographic algorithms were broken. In response, this last year national authorities and industry experts have intensified their calls to migrate to quantum-resistant cryptography.*

*Multiple quantum technologies have shown improved stability and scalability, providing promising pathways to a large-scale quantum computer. Experts now estimate that there is a 27% likelihood of a quantum computer breaking cryptography by 2034 with 1/3rd of experts assessing the likelihood at 50/50 by 2034. Furthermore, the US, Australia and several European nations have set timelines and guidance for the transition, with 2030 emerging as the probable pivotal date after which many organisations should not rely upon existing quantum-vulnerable asymmetric cryptography.*

*Organisations should be preparing* ***now*** *by assessing their risks and engaging their vendors to introduce quantum resistance ahead of the threat being realised, prioritising protection of long-lived sensitive data and the hardware security foundation. With general purpose quantum-resistant cryptography algorithms now standardised by NIST and being adopted internationally,* ***2025 is the first full year where most quantum-vulnerable implementations now have a viable migration path.*** *As a result, we expect to see protocols and products offering quantum resistance on a widespread scale.*

*Now is the time to ask vendors how they will be providing quantum-resistant protection.*

*Two significant false alarms of quantum breakthroughs sent jitters throughout the security community last year. Though these were effectively assessed, they serve to keep us alert to how damaging a real breakthrough could be on our digital infrastructure if we do not get ahead and prepare now.*

HP’s excellent state-of-the-race overview was heavily resourced with links to backup everything. I've included the link to their full article in the show notes.
https://threatresearch.ext.hp.com/protecting-cryptography-quantum-computers/

We really are in a time of significant change. Governments are tackling the tough problem of wanting to protect their citizen’s privacy while not wishing to allow criminals to evade responsibility for their crimes by abusing absolute privacy. The move from the physical to the cyber world has parents and guardians wishing to protect their children from online harms which means there’s no way getting around knowing at least something about who’s who. And on top of all this the fundamental technology that underlies any of our ability to do these things is strongly expected to collapse and be rendered useful once quantum computers – whose arrival now appears to be inevitable – are brought to bear.

We certainly are living through interesting times. And we’ll be back here next week on April Fool’s day to see what other “interesting” things have happened!