Security Now! #1002 - 11-26-24

Disconnected Experiences

This week on Security Now!

What's the new “nearest neighbor” attack and how do you defend against it? Let's Encrypt just turned 10. What changes has it wrought? Now the Coast Guard is worried about Chinese built ship-to-shore cranes. Pakistan becomes the first country to block Bluesky. There's a new way to get Git repos "swatted" and removed. Who's to blame for Palo Alto Networks' serious new 0-day vulnerabilities? If you have any of these six D-Link VPN routers, unplug them immediately!
It turns out that VPN apps are against Shariah Law. Who knew? The Return of Windows Recall. What are we learning now? How many of today's systems remain vulnerable to last year’s most popular exploits? We share and respond to a bunch of terrific feedback from our listeners. Then we ask: What are Microsoft's “Connected Experience” and why might you choose to disconnect from them?

What’s wrong with this picture?

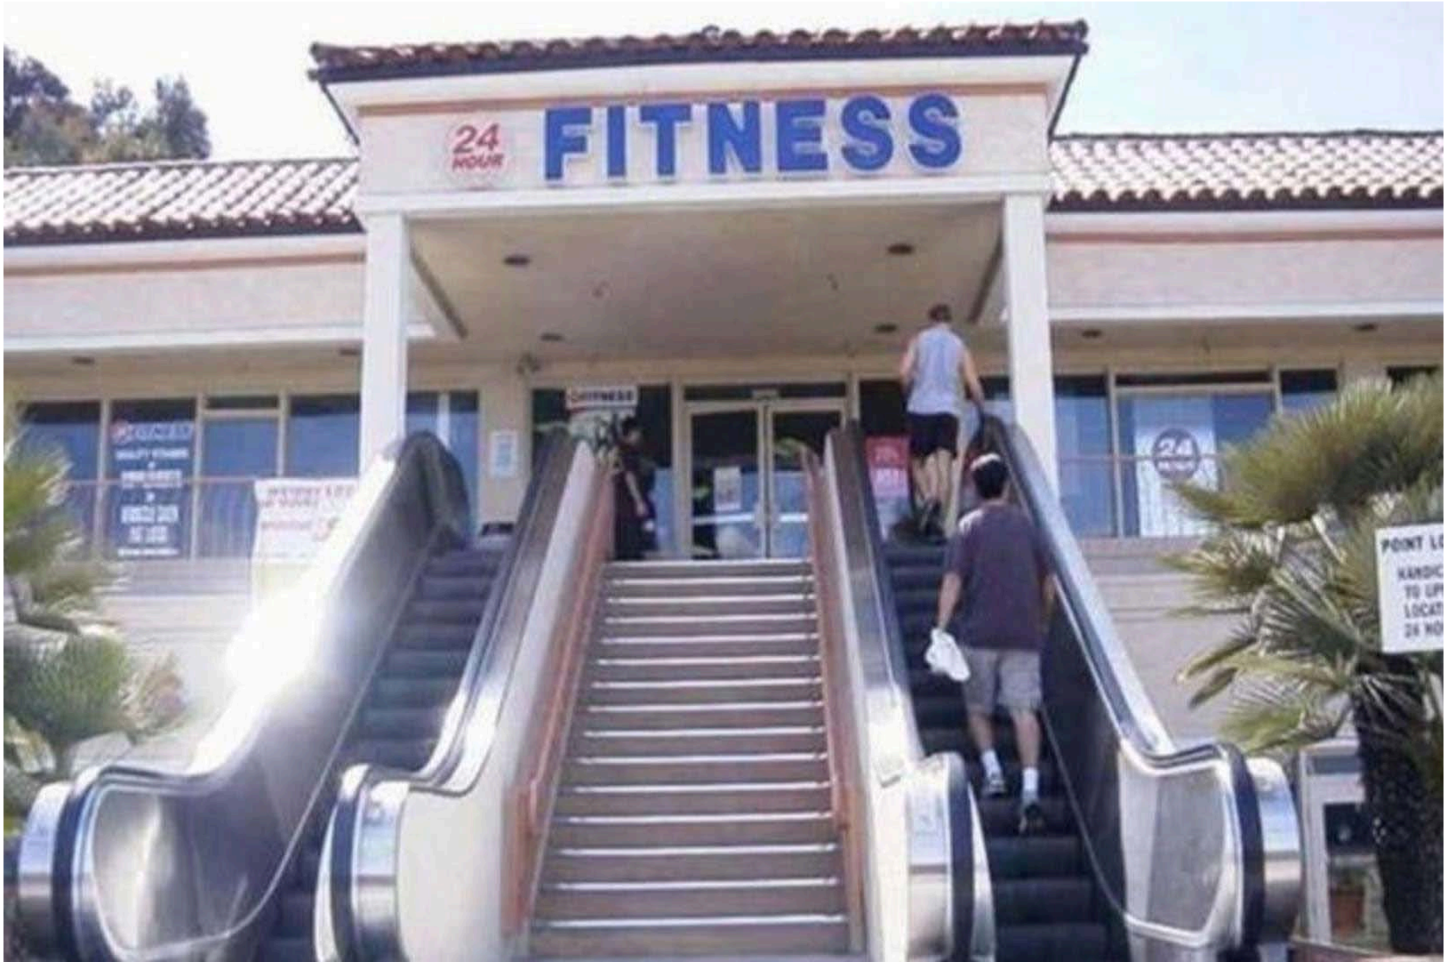

Security News

Volexity’s “Nearest Neighbor” Attack

Last Friday the 22nd, the security firm Volexity published the details of a somewhat astonishing and successful attack. Being several years old, predating Russia’s invasion of Ukraine, this story is not about a threat any of us will ever face – at least almost certainly not. I wanted to share it since it presents a perfect example of my “porosity” theory of security – where the security of today’s systems is best viewed as being porous to varying degrees. I believe the model of a porous system fits best because while the amount of effort an attacker may need to exert to obtain access to any specific system may vary, most systems can ultimately be breached by a sufficiently motivated and determined attacker. In other words, “absolute security” is more a concept than a reality today. Here’s how Volexity opened their disclosure. They wrote:

> *In early February 2022, notably just ahead of the Russian invasion of Ukraine, Volexity made a discovery that led to one of the most fascinating and complex incident investigations Volexity had ever worked. The investigation began when an alert from a custom detection signature Volexity had deployed at a customer site (we’ll refer to them as “Organization A”) indicated a threat actor had compromised a server on the customer’s network. While Volexity quickly investigated the threat activity, more questions were raised than answers due to a very motivated and skilled advanced persistent threat (APT) actor, who was using a novel attack vector Volexity had not previously encountered. At the end of the investigation, Volexity would tie the breach to a Russian threat actor it tracks as GruesomeLarch (publicly known as APT28, Forest Blizzard, Sofacy, Fancy Bear, among other names – in other words, Russians.) Volexity further determined that GruesomeLarch was actively targeting Organization A in order to collect data from individuals with expertise on and projects actively involving Ukraine.*

So what did Volexity’s investigation uncover? Strange as it might at first seem, despite being thousands of miles away in Russia, the well known APT28 group of Russian state sponsored actors breached an unnamed U.S. company by gaining access through its enterprise WiFi network. But wait... they’re thousands of miles away. How’s that possible? If I told you that the attack has been dubbed “the nearest neighbor attack” you’d start to get the idea. That’s right. APT28 pivoted to their ultimate target after first compromising an organization in a nearby building that was within WiFi range of their target.

APT28 has this level of expertise. They’re part of Russia's military unit 26165 in the General Staff Main Intelligence Directorate (the GRU) and they are known to have been conducting offensive cyber operations since at least 2004 – so for the past 20 years.

APT28 initially obtained the credentials to the target's enterprise WiFi network through password-spraying attacks targeting a victim's public-facing service. However, the presence of multi-factor authentication prevented the use of the credentials over the public web. Although connecting through the enterprise WiFi did not require MFA, as Volexity phrased it, “being thousands of miles away and an ocean apart from the victim” presented a problem.

So the hackers got creative and started looking at organizations in buildings nearby that could serve as a pivot to the target wireless network. The idea was to compromise another organization and search its network for a wired accessible device containing a wireless adapter. Such a device – a laptop, router or access point – would allow the hackers to use its wireless adapter to connect to the target's enterprise WiFi.   Volexity wrote this:

> *Volexity now determined the attacker was connecting to the network via wireless credentials they had brute-forced from an Internet-facing service. However, it was not clear where the attacker was physically that allowed them to connect to the Enterprise Wi-Fi to begin with. Further analysis of data available from Organization A's wireless controller showed which specific wireless access points the attacker was connecting to. After overlaying them on a map that had a layout of the building and specific floors, Volexity could see the attacker was connecting to the same three wireless access points that were in a conference room at the far end of the building near windows along the street. This gave Volexity the first evidence that the "call was not coming from inside the building." Could this be an attacker conducting a close access operation from the street outside? Nothing was ruled out, but Volexity was not too far off from discovering the real answer.*

What they discovered was that APT28 had compromised multiple organizations as part of this attack, daisy-chaining their connection using valid access credentials. Ultimately, they gained access to a device continuing a WiFi radio that was able to connect to three wireless access points near the windows of the victim's conference room. Then, using a remote desktop connection (RDP) from an unprivileged account, the threat actor was able to move laterally within the target network to search for systems of interest and to exfiltrate data. The attackers generally used "Living off the Land" techniques, relying on already present native Windows tools in order to minimize their footprint and thus reduce the chance of being detected.

Even with all of their research, Volexity was working from forensic data and was unable to trace the attacks back to their attacks. Attribution was impossible. But a Microsoft report last April provided the missing clues. Volexity saw clear overlap in indicators of compromise (IoCs) that clearly matched and pointed to the Russian advanced persistent threat group. Based on details in Microsoft's report, it's very likely that APT28 was able to escalate privileges before running critical payloads by exploiting a 0-day vulnerability CVE-2022-38028 that existed in the Windows Print Spooler service within the victim's network.

So our unsettling takeaway from this is that close-access operations that typically require proximity to the target, such as from an adjacent parking lot, can sometimes be conducted from great distances. In addition to making an otherwise impossible attack possible, this eliminates all the risk to the attacker of being physically identified or caught.

The other takeaway is that everything should be logged. The mantra should be "Log everything." It's crucial to appreciate that it's inherently impossible to know which logs will be needed. And nothing brings an investigation to a grinding halt more quickly than running up against "we don't have logs of that." Today's storage is so inexpensive that it's no longer a factor. Logs don't take up much space. They contain so much redundant formatting and repetitive data that they compress down to nothing. And they serve as a form of time machine that later allow forensics investigators to venture far back into the past to view what happened when and to retrace the previously unseen footsteps of unknown network users. And logs are not only useful for tracking Russians. Large corporations cannot be certain about the changing motivations of their own employees. So an IT culture of logging, and letting it be widely known that everything within an organization is being logged is a bit like planting a sign on the front lawn to let would-be burglars know that the premises is being monitored by such-and-such company. It can be an ounce of prevention.

Let's Encrypt turned 10.
Last Tuesday was the 10th anniversary of Let’s Encrypt and its statistics page shows that its certificates are now being used to encrypt the connections of 500 million domains:

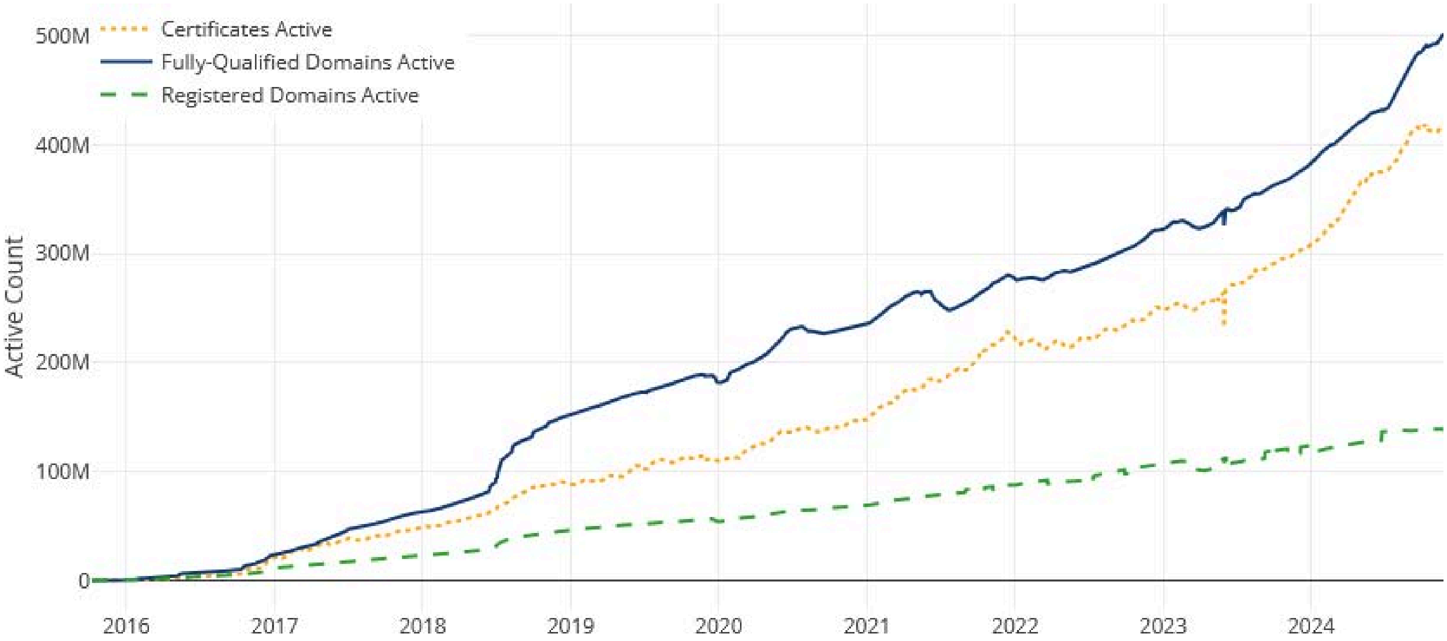

And the RATE of certificate issuance tells the story:

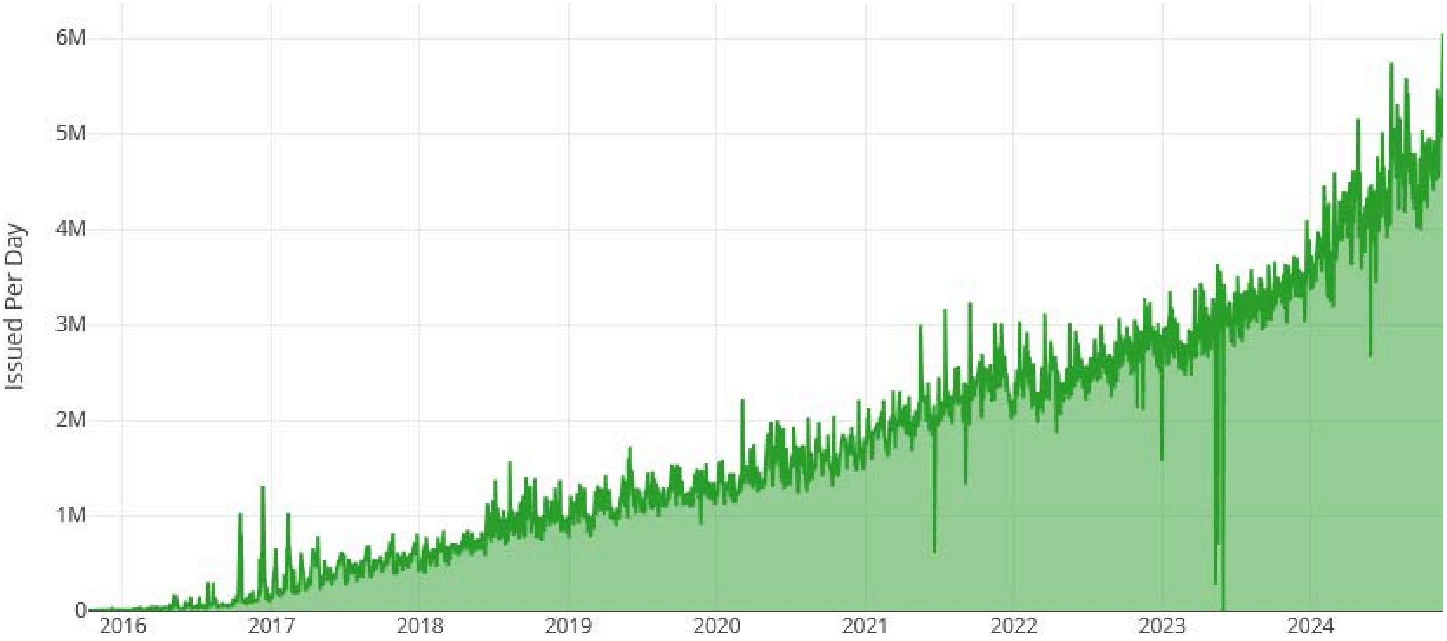

20 years ago when this podcast began, most websites used unencrypted and unauthenticated HTTP. Those sites which needed to obtain private and confidential information from their users, even if that was only their username and password to login, would typically switch to an HTTPS connection only during the transmission of that information. But as the world grew to become ever more dependent upon the Internet for everything, it became clear that the original “trust by default” model was not going to take us where we needed to go in the future.

The industry needed a future where the privacy provided by encryption could be available to everyone. The trouble was that encryption required certificates and certificate authorities had made a lucrative business out of verifying the identity of website owners and signing their certs which attested to that verification. And since performing this verification required a lot of work, certificates carrying those attestations were not free. The ISRG – the Internet Security Research Group – was formed to solve this problem. Two engineers from Mozilla, a guy from the EFF and one from the University of Michigan incorporated the ISRG and set to work solving this problem.

The Group decided that the inherently expensive and scaling-resistant verification of domain ownership could simply be bypassed in favor of reducing the test to anonymous domain control. And if that was done, web and DNS servers would be able to verify the domains they were serving and the entire process of certificate issuance and maintenance could be automated. Thus the ACME – Automated Certificate Management Environment – protocol was born. And today, half a billion domains later, by any measure this has all been a huge success.

Thanks to Let’s Encrypt, any website that wishes to can now have every connection encrypted for privacy. Have Let’s Encrypt’s free certificates been abused? Of course they have. That’s what happens on the Internet when anything is free. Look at email spam, and today’s social media. Both are an utter catastrophe because both are free. But this wasn’t the problem Let’s Encrypt was trying to solve or prevent. Their clearly stated goal was to offer equal opportunity privacy through encryption for all. Bad guys and phishing sites were every bit as welcome to have Let’s Encrypt certificates as anyone else. At least the communications of the people they were scamming would now be private and encrypted. And that really was all the ISRG intended to provide.

Coast Guard Warns of Continued Risks in Chinese Port Cranes

Last Wednesday’s report in GovInfoSecurity was titled “Coast Guard Warns of Continued Risks in Chinese Port Cranes”. This becomes an issue when it’s accompanied by the news that 80% of all heavy lift gantry cranes used to load and unload container ships at American ports were manufactured by ZPMC, a state-owned Chinese company.

The report explains that the U.S. Coast Guard is warning that Chinese-made ship-to-shore cranes come with “built-in vulnerabilities” enabling remote access and control. Consequently, the Coast Guard has begun urging operators across the country to adopt enhanced security protocols. In their notice, the Coast Guard wrote: “Additional measures are necessary to prevent a transportation security incident” and the Coast Guard cited “threat intelligence related to the PRC's interest in disrupting U.S. critical infrastructure.”

The notice instructs owners and operators of Chinese-made STS cranes – where STS stands for Ship-To-Shore – to obtain a copy of the official directive from their local Coast Guard officials, stating that the materials contain sensitive security information. A congressional report published in September warned a Chinese company with a major share of the global market of STS port cranes posed “significant cybersecurity and national security vulnerabilities” for the U.S.

According to the report, the Chinese state-owned company, known as ZPMC supplies 80% of all Ship-To-Shore cranes in the U.S. market and has significant involvement in militarizing the South

China Sea. Lawmakers warned that the company and its cranes could “serve as a Trojan horse” allowing Beijing to “exploit and manipulate U.S. maritime equipment and technology at their request.” What remains unclear is what measures the Coast Guard could implement to restrict the remote functionality of Ship-To-Shore cranes which are integral to port operations nationwide.

So we add this new example to Chinese-made DJI drones and Chinese-made security cameras which those in the U.S. have been blithely purchasing and plugging in everywhere for years. The answer to the question of what are we to do about these cranes is the same as for the DJI drones and cameras: In theory, we could purchase the hardware and independently source the firmware or software for these devices. But nothing prevents firmware buried deeply within the hardware from being similarly compromised. So the simple truth is, in any instance where we’ve seriously and firmly determined that we cannot trust the supplier of equipment, that equipment cannot be used anywhere its physical or cyber compromise might lead to other damage. Imagine if Beijing could do nothing more than cause 80% of all U.S. Ship-To-Shore port cranes to self-destruct. It would instantly and irreversibly cripple all major U.S. ports.

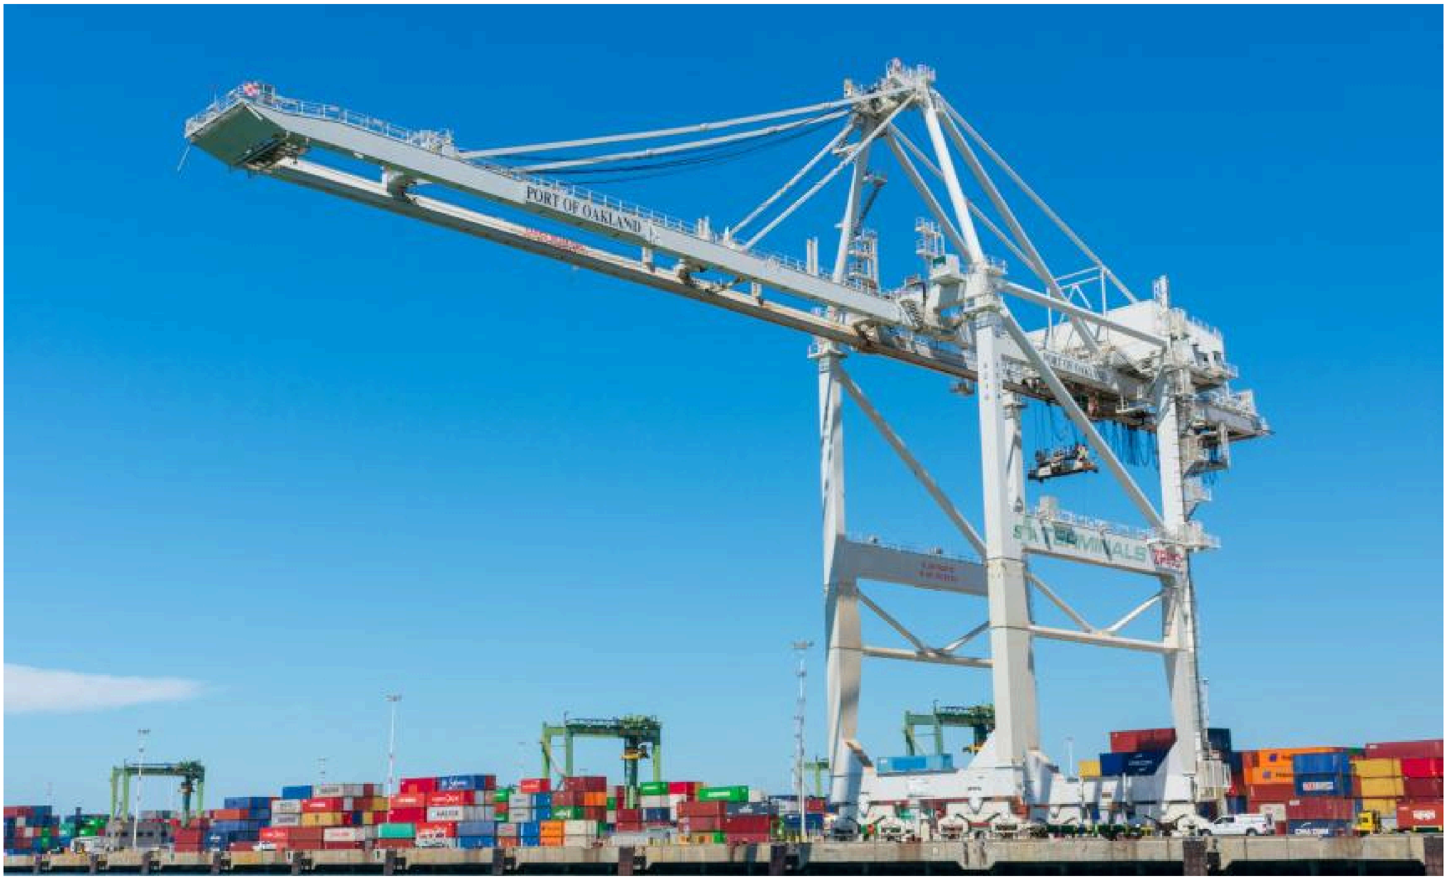

A ship-to-shore gantry crane at Oakland International Container Terminal built by ZPMC

And in the case of these cranes, that would be a crying shame. I have a photo of one of these beautiful machines in the show notes. Just look at it! It’s like something out of Star Wars. You definitely don’t want to have that thing walking in your direction. I want one. Except then, look at the itty-bitty size of the standardized containers next to it and the truck parked near its base, and you suddenly gain a sense of perspective about the size of this monster. What a beautiful machine. Pity we can’t trust it.

Pakistan is first to block BlueSky

After a phenomenal surge in new users, BlueSky has received its first country-level block, and ... the winner is Pakistan. For those who don't know, BlueSky was originally conceived as a project within Twitter by Jack Dorsey. It was designed to create an open, decentralized standard for social media and it was launched in 2021 as an independent entity. After that, BlueSky quickly evolved into a strong competitor to X, offering a more customizable and transparent user experience.

BlueSky's overall popularity has been soaring recently and in Pakistan specifically, this is being driven by increasing accessibility issues with X. Due to government restrictions and the growing need for a VPN to access X, many Pakistani users have turned to BlueSky as an alternative. Unfortunately, it appears that within Pakistan BlueSky is quickly hitting the same barriers.

I should mention that I've received Twitter (okay, 'X') DM's asking when I'll be moving to BlueSky. I'm not moving anywhere. For me, 'X' is being allowed to slowly fade. I'm still posting the weekly show notes to 'X' because I've been doing so for years and some of our listeners who hang out there continue to appreciate that. But a nicer presentation of today's show notes was emailed to more than thirteen and a quarter thousand of our listeners yesterday, and every one of those listeners is able to address email directly back to me at "securitynow@grc.com". **And** all of that works even for our listeners in Pakistan.

"Repo Swatting" Attack

Under the section "What will they think of next?" We now have what's being called "repo swatting attacks". REPO is, of course, short for "Repository" which is the unit of organization employed by GitHub and GitLab. So get a load of this: Threat actors have been abusing a hidden feature to cause GitHub and GitLab accounts to be taken down. The technique allows users to open issues against a targeted repo, upload a malicious file, and then abandon the issue without publishing it. On both GitHub and GitLab, the file remains attached to a victim's account. Then, the pesky threat actor reports the hidden, non-public file for breaking the service's Terms of Service which forces the repo to be removed for hosting malware. Apparently, this is just one more reason why we can't have nice things.

Palo Alto Networks 0-Day

A couple of weeks ago I touched on two recently announced 0-day flaws that had been discovered to affect Palo Alto Networks enterprise firewalls. That led to my quite predictable rant about the proven impossibility of protecting **any** form of remote management access to Internet-facing services. Even firms like Palo Alto Networks, whose business is security and security appliances, don't know how to do that yet.

In this case, to say that Palo Alto's internal architecture seems somewhat wanting would be an understatement. An analysis by WatchTowr Labs reveals that this vulnerable appliance is implemented in what they declare, with tongue in cheek, to be the "absolutely stellar PHP language" which is served by Apache fronted by an Nginx reverse proxy. They then note that the system implements its authentication layer using a PHP feature known as "auto_prepend_file" which prepends the file "uiEnvSetup.php" to anything PHP loads.

This is implemented by the line auto_prepend_file = uiEnvSetup.php in PHP's PHP.INI file which they preface by saying "Take a look at this gem of a hack in the php.ini file" – and I could not agree more. They introduce its use by noting: "We guess auto_prepend_file actually has legitimate use besides writing PHP exploits."

The bottom line is that this is all quite dispiriting. I don't know why I always imagined that Palo Alto Networks' would be doing things right. I suppose I wanted to give them the benefit of the doubt. The uiEnvSetup.php text file which provides front-end authentication by redirecting pre-authenticated access to the login page contains the comment: *"these are horrible hacks. This whole code should be removed and only made available to a few pages: main, debug, etc."* I couldn't agree more, and I would never say that Palo Alto Networks **deserves** to have been hit by these vulnerabilities, especially since it's their customers who will be taking the hit for this. But a design that is this slipshod can only be called ***"asking for it."*** It's unconscionable that this is the utter crap they're shipping. In order to see any of this, the WatchTowr guys needed to first jailbreak this Palo Alto appliance. This means that this extremely poor design is locked away out of sight so that it's only visible to intrepid researchers. But even if it cannot be seen, every Palo Alto customer remains reliant upon it.

We all know the rigid line I draw between bad policies, which are deliberate, and true mistakes which anyone could make. None of this is not an example of a mistake anyone could make. There are developers inside Palo Alto Networks who **know** this is what they are shipping. Those people should be looking for a new job far away from anything having to do with security.

So today we have the news from the The Shadowserver Foundation of evidence that at least 2,000 of those Palo Alto Networks firewalls have been compromised using those two recently disclosed 0-days. Once they have been compromised, the firewalls contain a PHP webshell which allows attackers to return later. The presence of this webshell is one indicator of compromise. The Shadowserver Foundation said that their number was a conservative estimate since it relies upon a limited set of IOCs released by Palo Alto Networks last week. To their credit, Palo Alto Networks had warned of a possible zero-day early this month and their communication throughout this has been stellar. So there's much to commend Palo Alto Networks about their response to this trouble. Unfortunately, this stands in stark contrast to whomever is developing their devices.

Hardware is not forever

A responsible security researcher going by the handle "delsploit", who reportedly answers email at "delsploit@gmail.com", has privately and responsibly disclosed their discovery of a terminally serious stack buffer overflow vulnerability across D-Link's past VPN routers. I characterize this as being terminally serious because this now-known-to-exist vulnerability allows unauthenticated users – also frequently referred to as "anyone anywhere" to remotely and at their whim execute their remote code on the victim's targeted D-Link VPN router.

The concerns are that D-Links announcement of this sobering reality last Monday contains a field for "Link to Public Disclosure" which is currently filled-in with the abbreviation "TBD" as in "To Be Determined" which strongly suggests that this "delsploit" character is being responsible with his or her knowledge and is giving D-Link some time to respond.

But there’s a problem with that: All six of these venerable (and vulnerable) D-Link VPN routers have gone well past their end of life. They are no longer being supported by D-Link and thus will not now, and not ever, be receiving updates to correct this most critical vulnerability. No CVS tracking designation has been assigned to track this vulnerability because it’s never going to be fixed, and if a CVS were to be assigned it would likely carry a distressing flashing red CVSS score of 9.8 or perhaps even the rarest of rare 10.0.

This vulnerability is as bad as they come because this otherwise lovely family of routers offers a standard SSL VPN which runs a simple web server at the standard HTTPS port 443:

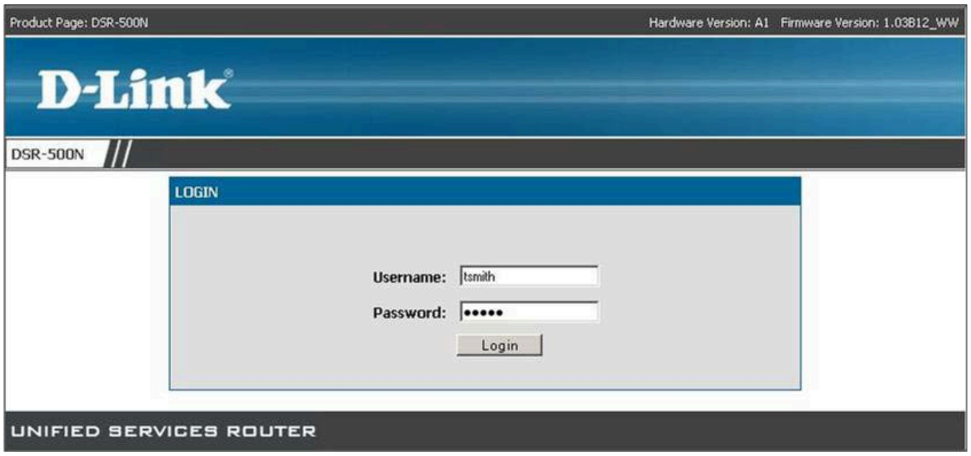

Product Page: DSR-500N

Hardware Version: A1 Firmware Version: 1.03B12_WW

Username: tsmith

Password: •••••

I have a screen shot of this in the show notes. From the standpoint of almost actively soliciting attackers this could not get any worse. The page that’s displayed to any device connecting to port 443 of an affected router prominently displays the device’s model number and both the hardware and firmware version numbers. This thing effectively shouts “please exploit me!” So “where they are” on the Internet will never be any mystery and I have no doubt that lists of their IP addresses have long ago been assembled.

Okay. So now everyone knows the situation. The two oldest affected models are the DSR-500N and -1000N both which went End-Of-Life nine years ago back in September of 2015. The more recent four VPN routers are the DSR-150, 150N, 250 and 250N. All four of those went End-Of-Life just a few months ago, in May of this year. But as the saying goes, “Close only counts in horseshoes and handgrenades” meaning, in this case, that end of life is end of life and that D-Link formally states in their disclosure that these now known to be seriously vulnerable D-Link VPN routers will never receive updates.

Longtime listeners of this podcast know what will come next, as sure as the sun rises every morning. Many tens of thousands of these devices are currently sitting on the public Internet. I haven’t seen an exact count but I’m sure that either Shodan or Censys would have that number – and be able to provide their IP addresses – since every one of them proudly presents its logon web page to any passer by. There’s been no public disclosure of the details of the vulnerability that “delsploit” found. But D-Link has confirmed it. And at some point “delsploit” is going to want their day in the sun and bragging rights about this vulnerability. So it’s going to be published. And no one can really fault “delsploit” for eventually disclosing the vulnerability they discovered because that’s the way the game is played these days: You wait long enough to give the impacted parties a reasonable amount of time to respond, and after that, no matter whether or not they have, and regardless of the consequences, the entire hacking elite is then informed of exactly how to bypass the Internet-facing authentication which protects the tens of thousands of networks behind every one of these VPN routers.

There's nothing any of us can do other than protect ourselves and those we have responsibility for and care for. Make absolutely double-damn certain that nowhere within your spheres of influence do any of these six D-Link VPN routers currently exist. Because we all know exactly what's going to happen next.

In their disclosure, D-Link ineffectually recommended that this hardware should be replaced. We know that most of the owners of these devices will never receive any sort of notice of this; and probably wouldn't pay it the attention it deserves even if they did. We're all being so inundated by all of our software being constantly updated that it's easy to become numb to it. But if anyone is in the market for a replacement I would stay well clear of D-Link. They have a long and still-growing history of very serious remotely exploitable vulnerabilities being discovered after the fact in their past end-of-life products. This happened earlier this month with 66,000 of D-Link's Internet-connected NAS devices. Their response was effectively: "We're sorry, we don't make NASs any longer. And even if we did, those 66,000 Internet connected remotely exploitable NAS devices we once made are now past end-of-life, so it wouldn't matter anyway." It's true that hardware is not forever. And that it would not be unreasonable to expect an aging NAS or router that's past its end-of-life to be rotated out of service in favor of something new. But we all know that this often doesn't happen. Given their track record, I would be disinclined to give D-Link any more commercial support. If you really like the brand, I get it, it's truly nice looking hardware. But you should be aware that "end of life" or "end of support" probably means "end of secure service life" after which point a device should be rotated out of service. And if you have an existing inventory of D-Link devices you should be very certain to have a current subscription to their security bulletins and other notifications.

Use of VPNs is against Shariah Law

"Sharia" is a religious law that governs the lives of Muslims, based on the teachings of Islam and the Quran. So while we're on the topic of Pakistan being unhappy with pretty much all things Internet, I should note that Pakistan's religious advisory board recently ruled that the use of VPN apps is against Shariah Law, apparently because, you know, "Shariah Law" is pretty much anything we want it to be.

The Council of Islamic Ideology said that VPN technology was being used in Pakistan to access content prohibited according to Islamic principles or forbidden by law, including *"immoral and porn websites or websites that spread anarchy through disinformation."* The council's chairman said that the use of VPNs falls under 'abetting in sin.'

This gave me pause to wonder, whether they might be inclined to change their minds if they were able to get a great deal on some used D-Link VPN routers? What do ya think?

The Return of Recall

Last Friday, the Windows Insiders Blog announced the return of Recall to Windows 11. They wrote: *"Hello Windows Insiders, today we are releasing Windows 11 Insider Preview Build 26120.2415 (KB5046723) to the Dev Channel. With this update, we welcome Windows Insiders*

with Snapdragon-powered Copilot+ PCs to join the Dev Channel to try out Recall (Preview) with Click to Do (Preview).”
https://blogs.windows.com/windows-insider/2024/11/22/previewing-recall-with-click-to-do-on-copilot-pcs-with-windows-insiders-in-the-dev-channel/

I have a link to the lengthy roll-out text in the show notes for anyone who wants more. Suffice to say that Microsoft has done exactly what they had promised to do. The setup experience of course promotes Recall as a wonderful and fully secure feature. It’s unclear from the few screenshots Microsoft provided what the user’s decision tree looks like and how readily the user is able to decline to receive the “Recall experience”. But presumably, after all the backlash Microsoft previously received and their commitment to disable Recall until and unless its user explicitly enabled it, that’s what they have done. I do know from reporting that Recall can mostly be removed from Windows through the “Turn Windows feature on or off” dialog. One security researcher noted that a few Recall-related DLL’s remain under \Windows\SystemApps\, specifically a file named MicrosoftWindows.Client.AIX. But this researcher noted that the core functionality is removed.

A few items of note from the blog posting were, quote:

> *Recall (Preview) will begin to rollout on Snapdragon-powered Copilot+ PCs, with support for AMD and Intel-powered Copilot+ PCs coming soon. As we gradually roll out Recall in preview, Recall is supported on select languages including Chinese (simplified), English, French, German, Japanese, and Spanish. Content-based and storage limitations apply. Recall is not yet available in all regions, with expanded availability coming over time.*

There were anecdotal reports of researchers being able to get the first shot at Recall running on PCs without any fancy AI GPU support. So it might be that Recall will be made more widely available over time. This might also mean that for now, no one without a Copilot+ PC will need to worry about removing it since it may never be present.

Also of interest in the posting, Microsoft wrote for Enterprise customers:

> *As announced at Ignite, for our enterprise customers, Recall is removed by default on PCs managed by an IT administrator for work or school, as well as Enterprise versions of Windows 11. IT administrators fully control the availability of Recall within their organization. Employees must choose to opt-in to saving snapshots and enroll their face or fingerprint with Windows Hello for snapshots to be saved. Only the signed-in user can access and decrypt Recall data, so although enterprises cannot access employee Recall data, they can prevent Recall from being used altogether and prevent any saving of specific apps or sites.*

Just for the record, Microsoft is also previewing a Recall feature they call “Click to Do”, and they write:

> *With Click to Do in Recall, you can get more done with snapshots and improve your productivity and creativity. Click to Do recognizes text and images in snapshots and offers AI powered actions you can take on these, saving you time by helping complete tasks inline, and/or quickly getting you to the app that can best complete the job for you.*

They then show that the user has been able to Mark and highlight to Select text in an image on a Recalled screenshot. Once selected, a context menu offers commands to:

- **Copy**: Easily copy text to your clipboard.
- **Open with**: Open the selected text with your preferred application.
- **Search the web**: Quickly search the web for the selected text.
- **Open website:** Open any URL you recognize on screen in your preferred browser
- **Send email:** Send email to the email address recognized on screen in your preferred email app

And if the user right-clicked on a Recalled image, the context menu commands would then be:

- **Copy**: Copy the image to your clipboard.
- **Save as**: Save the image to your desired location.
- **Share**: Share the image with others.
- **Open with**: Open the image with your preferred application.
- **Visual search with Bing**: Perform a visual search and surface relevant contents using Bing.
- **Blur background with Photos**: Blur the background of the image using Photos app.
- **Erase objects with Photos**: Erase unwanted objects from the image using Photos app.
- **Remove background with Paint**: Remove the background of the image using Paint app.

Microsoft explains:

> *In this update Click to Do only works within the Recall experience. In a future update, you'll be able to effortlessly engage with Click to Do by simply pressing Windows logo key + mouse click, Windows logo key + Q, through the snipping tool menu and Print Screen, or searching "Click to Do" through Windows Search Box. These methods will make it easier than ever to take immediate action on whatever catches your eye on-screen. We're also working on introducing more intelligent text actions to enhance your experience even further.*
>
> *Just like with Recall noted above, Click to Do (Preview) is available only on Snapdragon-powered Copilot+ PCs. Support for Intel and AMD-powered Copilot+ PCs is coming soon.*

We need to (and can) do better!

I was talking earlier about the fact that we absolutely know that very very few of the now known to be vulnerable D-Link VPN routers will be removed from the Internet as a result of D-Link's announcement of their serious vulnerability. How do we know? Well, CISA maintains a list of the most exploited security vulnerabilities by year.

We know that at least 60 known threat actors exploited vulnerabilities from CISA's list of the most exploited bugs last year. And we have details. According to the security firm VulnCheck, the North Korean group "Silent Chollima" was the most active in this regard. They targeted 9 out of 15 CVE's from CISA's list. China and Russia's groups were the most active among the 60 known threat actors with China sponsoring 15 of those 60 and Russia supporting 9.

And here's the most distressing news that gets back to why we know that few of those D-Link routers will be removed from service: VulnCheck reports that over 400,000 systems that are currently online at this moment are vulnerable to attacks using one of last year's most popular vulnerabilities. As an industry we really do need to do better.

Closing The Loop

Thomas

> *On a recent episode, you mentioned a device that acts like a bluetooth keyboard and connects via a dongle between a phone (or other bluetooth device) and a computer (or basically anything you could plug a usb keyboard into. It sounds to me like an input stick. http://inputstick.com/ (not https?) a device that I used frequently as a hardware tech when replacing HP motherboards. After you replaced the motherboard, you had to enter a setup "command?" "string?" that was about 30 characters long and case sensitive. Since it was entered before/during bios, you couldn't copy it into the field from the web. It was a nightmare. But with the input stick, you could go to HP's website on the phone, copy the string, paste it into input sticks software, and send it / input it correctly the first time. Been a while since I've done that. Now it mostly works as a volume control to turn my computer down when I'm going to sleep. Still one of my favorite toys though. Even though I'm no longer in the biz, I still keep up with the news via SecurityNow. Thomas*

Thomas is 100% correct. The gizmo another listener mentioned, which I immediately purchased since it looks clever and interesting. It was $39 US plus shipping from Poland. I'll report again once I've had the chance to play with it. Its creator appears to have done quite a lot with the capability, with the dongle able to simulate both a keyboard and a mouse, allowing simulation of multi-media control keystrokes, macros, etc.

I'm constantly annoyed that despite my decades long loyalty to all things Apple for everything other than PCs, Macs offer integration features that Apple refuses to bring to Windows. On the other hand, if what they were to bring to Windows resembled that abomination known as iTunes we might be better off without it.

In any event, thanks, Thomas.

Gino Guidi who signed his note "The Network Ninja" earns his title. He wrote:

> *Steve, Was listening to the episode where you had a user ask about how to capture the C2 traffic when it is using a hard coded IP. The solution you offered would absolutely work. I think the more elegant solution would be to just NAT the destination. I am not entirely familiar with pfSense or OPNSense as I use Untangle and Palo Alto at home. However, if you have FW software that supports it you could create a NAT rule that changes the destination from the hard coded IP to a host of your choice. You won't even need additional interfaces. If you configure the rule correctly it will re-NAT it back for return traffic. The malware will have no idea that it isn't actually talking to that IP. The additional advantage is that you wouldn't have to change the IP or add additional IPs onto the machine you are sending the C2 traffic to. You could easily create as many of those NAT rules as you want, which I think would make it more robust long term. I appreciate the podcast and hope to be listening for another 1000 episodes. Hope this suggestion makes sense.*

Gino: Given that a router's firewall software supports it, it's an absolutely brilliant solution that's clearly superior to the more complex approach I proposed. I like it a lot.

So let's think this through... As I understand it, it would require routing software that's able to perform NAT translation for packets traversing the router's internal LAN interface. That's different from typical consumer router NAT which is generally applied to outbound packets crossing the router's WAN interface. So this would definitely require some 3rd-party routing software.

Applying NAT to the internal interface would cause any packet sent from any machine on the LAN, such as the malware-infected machine, which is addressed to a specific external public IP to have its destination IP changed to another host machine on the LAN – the one that's serving as the command and control server. The packet's source IP would remain unchanged, the IP of the malware-infected machine.

So, on its way out from the malware-infected machine, the outbound packet crosses the LAN's selective NAT translation, which would give it a local destination LAN IP address. This would cause the router to send it back out the same LAN interface, addressed to the command and control server. And since that packet arriving at the C2 server would still be carrying the local source IP of the malware-infected machine, the spoofed C2 server would return its replies directly to the malware-infected server. It's an elegant solution and I can't see any reason why it would not work. Nice going "Network Ninja!"

Abhi Rau, driving his kids to school in Charlotte, North Carolina, wrote:

> *Hi Steve, I have been listening for the last 12 years. Your podcast has been a constant on my drive to work and dropping my kids to/from school. My kids have grown up listening to your voice and are more security conscious because of you. So thank you!*
>
> *In your last show (episode #1001) you mentioned Cloudflare Tunnel as an option for accessing home networks. One main clarification I would like to make, which you didn't mention, is that although a Cloudflare Tunnel is simple to setup and use, it does not provide true end-to-end encryption. While it encrypts traffic between your origin server and Cloudflare's network, Cloudflare can decrypt and inspect the data in transit as it terminates the TLS connection at its edge network, meaning it is not fully encrypted from start to finish.*
>
> *For true end-to-end encryption an overlay network like Tailscale can be used. For more detailed comparison: https://tailscale.com/compare/cloudflare-access. I looked into Cloudflare Tunnel to access my self hosted Bitwarden running on my home Synology NAS but I decided to use Tailscale instead for this reason. Love the show. To 2000 and beyond!*

Abhi provided a link to Tailscale's Tailscale-vs-Cloudflare-Tunnel side-by-side feature comparison which I have in the show notes for anyone who's interested. And I tend to agree with Abhi's feelings. I think that the best way to think of it is that these two solutions have some overlap which allows either one to solve the remote access problem, but they are also very different. Cloudflare Tunnel has a large range of features that go far beyond what's needed to remote access to a user's LAN. It's really aimed at secure remote access to servers. And an overlay network's end-to-end encryption is really what we want for remote network access.

Stephen Clowater reminds us of an even simpler solution:

> *Hey Steve, Congrats on hitting 1000+ episodes! Thanks for all the thoughtful content you've shared. I wanted to share an observation about remote access in homelabs, having tried Cloudflare Tunnels and various VPN clients. For those who don't need the features of an overlay like Tailscale,* ***WireGuard*** *is worth considering. It offers simple, lightweight Layer 3 connectivity, modern elliptic curve crypto, and straightforward setup. While Tailscale builds on WireGuard for robust overlay features, a standalone deployment keeps things minimal and widely supported across platforms like Linux, pfSense, and OPNsense.*
>
> *What has kept me using WireGuard is how it handles iOS sleep cycles, ensuring apps can reliably access data when waking from sleep. VPNs like OpenVPN, CF WARP, and IKEv2 often struggle with app-level connection failures because their clients may not wake up properly in the selective sleep process iOS has or renegotiate stale connections before a TCP timeout. WireGuard's small kernel footprint and fast connection negotiation allow it to reconnect on demand without timeouts.*
>
> *I started using WireGuard in 2020-2021 while setting up a self-hosted email server. I needed a reliable way to fetch mail on my phone while keeping port exposure to a minimum. Since then, it's become a core part of my setup, enabling reliable email fetch cycles, isolating Ubiquiti cameras, and syncing files via Syncthing on my phone. Just thought I'd share in case it's helpful to anyone exploring options. Best, Another Steve*

I'm really glad Stephen reminded us of the many benefits of just plain old Wireguard. We originally discussed Wireguard – the replacement for OpenVPN which had grown very old and stale – back when it first appeared on the scene about five years ago. In episode 744 I first talked about Wireguard after meeting and being very impressed by the founders of the Mullvad VPN service and learning that they were already adopting Wireguard. And recall that some time later, Linus Torvalds incorporated Wireguard natively into the Linux kernel.

The only downside to running, for example, Wireguard on a pfSense or OPNsense router is that the first thing you need to do is open a static port through the router's WAN interface to the Wireguard service. And from then on that port is open, facing the outside world, and you're relying upon Wireguard not to have any critical vulnerability that would allow an authentication bypass. If you're okay with that then Wireguard is likely the lightest weight and most secure solution available. And I loved what Stephen shared about its compatibility with iOS.

But running with a statically open port – which is never required when using any of the overlay networks – would tend to bend me away from Wireguard, much as I would otherwise love to be able to use it. What I would consider as an option would be adding some sort of port-knocking solution that would allow a remote IP to be authenticated so that that IP and only that IP could then connect to the Wireguard VPN running in the homebase router. Since an ICMP ping packet can contain plenty of payload, a simple and secure challenge/response that incorporates the endpoint IP addresses and some crypt would do the trick. If only there were more hours in the day!

Enrico Ng gave his note the subject: “EP989: backdoor or incompetence”

> *Happy 1000. I'm still a bit behind. I'm listening to EP989 where you talked about the chinese RFID badge chip that was found to have a backdoor.*
>
> *We've heard plenty of reports about vulnerabilities found where the manufacturer left some debugging credentials in. We've also heard lots of reports about backdoors in products. I'm curious, in general, how does one determine if something is a backdoor or incompetence? How can the researcher infer intent? Perhaps an internal company memo gets leaked that shows it was on purpose. It is still hard to tell if this was mandated by the government unless top secret documents get leaked. Is it just based on the country that manufactured the device and whether they are friendly to the US?*
>
> *I also heard about the guy that has gone back and started listening to your podcast from episode 1. I've wanted to do this, too. However, I'm already over 10 episodes behind so I'd just fall even further behind. I only listen to podcasts while driving. Maybe I need to plan some long road trips!*

I think that Enrico makes a very valid point. Controversy is inherent when attempting to ascribe intent. The question of the Windows Metafile Escape I noted last week is another perfect example. Why was it there? Why had it been faithfully copied and reimplemented through many editions of Windows, even jumping from Windows 3, 95, 98 and ME over to the brand new Windows NT... if it was an accident? The original intent of its designers has been lost to history and we’ll probably never know.

And remember about ten years ago when Cisco kept “discovering” hidden backdoor credentials in their appliances, one after another, month after month? (I have “discovering” in quotes because these were their own systems. How difficult could it be to “discover” a undocumented login account in software they wrote and for which they have the source code?)

Anyway, since Cisco is not evil, and never was, and since they were confessing over and over to what they kept finding, I think that is a case of poor judgement and the changing times. Twenty years ago, just as it may have been acceptable to design an escape hatch into Windows Metafiles, it may have been acceptable for developers to just leave their development accounts in Cisco appliance firmware. Back then it may have been no big deal. But as we’ve seen, times change our expectations.

My feeling is that in nearly all cases it’s just a mistake. For one thing, no clever developer would implement something that was meant to remain a secret by leaving a username and password in the firmware. That’s way too obvious. If someone told any competent developer to design-in a backdoor, it would be far more well hidden. For example, it would be necessary to first bounce an ICMP PING packet off the device with a particular payload length. This would leave an insignificant trace. Then it would be done again with a different specific length. And that pair of events would prime the device to then accept anything originating from the same source IP without requiring any authentication. Or something like that. My point is, nothing as dumb and obvious as leaving a username and password burned into the firmware. There are an infinite number of ways to bury a true backdoor in today’s insanely complex systems.

David (in the USA) wrote:

> *Hello Steve, I'm a long-time listener, but haven't reached out before. I credit you in large part for my career in infosec. I was unable to get formal education in the field, so I "self-taught" using resources including your podcast. It has been many years since I started my first job in the field, but I still listen regularly, and learn a lot. Thank you for all of your efforts!*
>
> *I'm sure this is an edge-case, but regarding your remarks about SoHo routers in SN-995, I was recently treated to an experience with a new Nokia (they still exist!) SoHo router/access point. I changed ISPs, and they provided one for "free," with a wifi SSID ready to use. They came out and installed it for me, and plugged what they thought was "my computer" into it (as if I only have one...haha!).*
>
> *After they left, I plugged my entire home infrastructure into their router. As a result of your recommendations some years ago, my main firewall pfSense running on a Protectli unit.*
>
> *I didn't bother to reconfigure the new Nokia box for a couple days, because I didn't consider it an important layer of security. However, I finally got around to logging into it, and was stunned by what I found. For some unfathomable reason, the firewall was set to a "light" filtering mode. Apparently it had a short, self-described "non-disruptive" block list it was using, to blacklist certain things. However, it was **not** performing NAT services for the ethernet!*
>
> *It was in a pass-through mode by default, giving my public IP address to my pfSense firewall behind it. There was an option on the Nokia device, to enable NAT, but it was disabled. While I would like to think that perhaps it detected the firewall behind it and switched itself off, I somehow doubt it is that smart. If I was a typical user, whatever I plugged into that Ethernet port would have been immediately exposed. The WiFi did seem to be using NAT, so perhaps they thought that was good enough for most users.*

The thing that occurred to me after thinking about what David wrote was that I'll bet almost no typical Internet user ever plugs anything into their router's wired Ethernet ports any longer. I know that many of us who listen to this podcast doubtless do. But we're far from typical Internet users. WiiFi really has overtaken wired Ethernet. That's the only way I can think to explain what David experienced.

Miscellany

The reMarkable Pro:

I have a bit of space at the bottom of this final page before we switch to today's main topic. So I wanted to answer the many questions I've received from our listeners who've taken note of the reMarkable Pro box on the bookshelf behind me, wanting to know what I think of it. I very much wanted to love it, but I don't. I wanted to like its support for color, its slightly higher pixel density, its larger size and its reputed higher stylus tracking rate. But I don't. It's support for color feels like it's not ready for prime time. The display goes through all sorts of conniptions when using color. It's clearly not easy to pull off color and I don't think it was worth it. Also, the darn thing is HEAVY. Really heavy. And its stylus now requires charging. By comparison, its predecessor, the reMarkable 2, I **really** love. I wish I could get the cool cover of the Pro which securely captures the stylus. But at least for the time being that appears to Pro only.

Disconnected Experiences

Microsoft silently enabled AI training for Word and Excel
The way things are going, it appears I'll be needing to set up a "Sacrificial Lamb" PC running the current, which is to say the latest, Windows. The last thing I would use for myself would be such a machine, because Microsoft really does appear to be pushing well past the limits of what is acceptable practice – "Windows Recall" is a perfect case in point. But if this podcast is going to continue to be as relevant as it has been in the past, it's becoming clear to me that I'll need to have a machine that's running what the rest of the unwashed masses are running – which is to say, the latest version of Windows.

There was a time when creating a sacrificial lamb PC meant exposing the machine to the Internet without protection. As we know, the half-life of such machines was best measured in seconds – and not many of those. But the way the Windows desktop environment has been evolving, today the creation of a sacrificial lamb PC means just exposing a machine to Microsoft.

The need for such a PC became clear when I encountered the news that Microsoft has silently enabled the use of its user's Microsoft Office Word and Excel document content for training its AI models. Rather than being straightforward and calling this something like, oh I don't know, how about "AI Training", they obscure it behind the title "Microsoft Connected Experiences". How the hell would anyone ever know what that means or what's going on? Connected Experiences?
And this is my point. This is what Windows has become. At the moment I'm reporting this blind because no way do I want to have anything to do with Windows 11 and whatever they're already doing today, let alone what they apparently believe they have the unlimited license to do in the future. Oh, but it's all free now, right? Nevertheless, if I'm going to be reporting on events which affect this industry I will need to be seeing these things for myself.

In Microsoft's documentation for their so-called *"Connected Experiences"*, under the topic *"Connected experiences that analyze your content"* they write: *"Connected experiences that analyze your content are experiences that use your Office content to provide you with design recommendations, editing suggestions, data insights, and similar features."* The key phrases there are *"analyze your content"* and *"connected"* – but connected to what and to where? That appears to mean what the reporting on this states, which is that the connection is to some AI which is doing the analyzing and being trained against Windows' users' Office document data. Now add to this the fact that it has reportedly been enabled by default. Because... of course it has.

It seems clear that just as a great many people are made uncomfortable by the idea of having Windows Recall silently collecting and analyzing everything they do on their computers, some Windows users may not be interested in having Microsoft's AI being trained on the content of their otherwise private Word and Excel Office documents.

First I'll note where this *"Connected Experiences"* setting is located since they clearly want their Windows users to have ready access to this potentially significant privacy setting. So under *"File"* in an Office application, choose *"Options"*. Under *"Options"* go to *"Trust Center"*. In the *"Trust Center"* select *"Trust Center Settings"*. There you'll find *"Privacy Options"* which you need to select in order to get to the *“Privacy Settings”* and on the privacy *“Privacy Settings”* page there’s a section for *“Optional Connected Experiences”* where you should find a checkbox labeled *“Turn on optional connected experiences”* which, all regular users will reportedly find has been thoughtfully enabled by default. Users whose machines or Microsoft accounts are managed by their organization may not have these options showing.

And Microsoft appears to confirm this on their own website, where under the topic *“Choose whether these connected experiences are available to use”* they write:

> *You can choose whether certain types of connected experiences, such as connected experiences that download online content, are available to use. How you make that choice depends on whether you're signed into Office with a Microsoft account, such as a personal outlook.com email address, or with a work or school account.*
>
> *If you're signed in with a Microsoft account, open an Office app, such as Word, and go to File > Account > Account Privacy > Manage Settings.*
>
> [Note that this is a different path from what I had just shared from the reporting of this. I don’t know which is correct. Perhaps they both are. But this is an example of why I’m going to need to be able to see these things for myself. Anyway, Microsoft continues...]
>
> *Under the Connected experiences section, you can choose whether certain types of connected experiences,* ***such as experiences that analyze your content,*** *are available to use.* ***If you don't go to Manage Settings, all connected experiences are available to you.***

So there it is: “connected experiences that analyze your content” and “If you don't go to Manage Settings, all connected experiences are available to you.” In other words, all connected experiences are enabled by default–which in turn means that all of your content will be analyzed.

What’s apparent nowhere is that “connected experiences” is a euphemism for “we’re going to share all of your Office documents to train an AI in the cloud in order to make Office smarter for you” – and, of course, for ourselves.

Talking about content retention they write: *“Most connected experiences don’t retain your content after performing their function to help you accomplish a task, but there are a few exceptions. In those cases, Microsoft retains the content for as long as your account exists and it’s used to support, personalize, or improve that connected experience.”*

As I write this, part of me wonders whether I’m just being an old curmudgeon? Why not just enjoy all of the many benefits of having Microsoft watching everything I do on my PC, thus allowing me to scroll back in time and ask questions about things I did in prior years. And sending my document content to the cloud to train their AIs so that it can provide me with more relevant stories on Edge’s home page, more relevant search results in Bing, and more relevant advertising on my Windows Start menu? I’m not being facetious when I say that many Windows users might want all of that. Just as they may have enjoyed having Candy Crush Soda Saga or whatever all that flippy-tile nonsense is under Windows 10 along with XBox crap that refuses to be removed. I’ve never owned an XBox but it has taken up residence on my Start menu nevertheless.

It seems clear that an alternative view of Windows is apparently an all encompassing deeply connected entertainment portal that also has some productivity applications. And, really, that's fine. It's just not for me. I mentioned a while back about the eventual move I would make to Windows 10 when I finally decide to retire this Windows 7 machine that's still working great. I was briefly thinking that a server edition might allow me to avoid some of this commercial crap – before I remembered that I had tried that years ago when I wanted my desktop to be running the identical code as GRC's servers. But I had encountered many instances of desktop software refusing to install on server editions. Some of our listeners have since suggested that I take a look at the enterprise editions of Windows 10, explaining that unlike even the Professional editions, the enterprise editions are also free of XBox and other unwanted nonsense.

And as I was digging around in Microsoft's documentation, I was encountering all of the places where Microsoft has been and is installing AI. Microsoft is essentially AI-izing every nook and cranny of Windows 11 and their Office suite. I have no doubt that a memo went out a year or two ago stating that AI was coming, that it was the future, and that once it had arrived it was here to stay. Therefore, every single product manager and product planning team within Microsoft was hereby being tasked with figuring out anything and everything that adding AI to their offerings could do, and to then get going on implementing all of that immediately.

What that will turn Windows into, I have no idea. I know that it won't be any machine that I'm sitting in front of while I produce these weekly Security Now! podcasts, nor while I'm working on code for the DNS Benchmark, the "Beyond Recall" product or SpinRites 7, 8 and 9. But it's also clear that I need to stay in touch with the frontier or, as it may more appropriately be termed... "The Bleeding Edge".

For now, I wanted to be certain that those listeners of ours, and I know there are many of them, who may dislike the idea of Microsoft sharing of their Office content with their AIs in the cloud, while acknowledging that this is being done by default and that in many cases the data is being retained indefinitely, would be informed of this new behavior and would know that they have the option of deliberately disconnecting their Windows experiences ... from Microsoft.

And finally, for those listeners who celebrate Thanksgiving, I know that Leo and I and all of the TWiT crew wish you the best holiday with its opportunity to spend precious time with family and friends.

We'll be back next week, in December, for more!