Security Now! #931 - 07-18-23

Satellite Insecurity, Part 1

This week on Security Now!

What did Kaspersky have to say about last Tuesday's Microsoft patch event, and what security consequences does it have for all non-subscribing Microsoft Office users? What was inevitably going to happen once the power of Large Language Model generative AI became widely appreciated and available? What does it mean that Microsoft just revoked more than 100 malicious Windows drivers? What two new well-known companies have been added to Clop's MOVEit file transfer victim list? What does Dun & Bradstreet have to do with Android Apps? Where in the world can you use Meta's new Threads service, and where not? And what's a side effect of bitcoin addresses looking like gibberish? And after we examine those questions, cover some miscellany and user feedback, we're going to turn our attention to the heavens in recollection of those famous words of Henny Penny.

Insecure Parking Spaces – Lock Your Car!

Security News

Kaspersky on Microsoft's Patch Tuesday

Kaspersky being Kaspersky, a very technologically savvy security firm, they had an interesting take on last Tuesday's monthly Microsoft patch event. It was the heading on their posting that drew me in. They titled their posting *"Band-aid on a… corpse: Microsoft patches IE — again"* and their sub head was *"July Microsoft Patch Tuesday: a collection of exploited vulnerabilities."* So this is all definitely worth sharing as we look back at the past week. They wrote:

*The Microsoft July patch collection has turned out to be quite surprising. First, they're once again fixing apparently dead Internet Explorer. Second, as many as **6** of the vulnerabilities are already being actively exploited by attackers. Third, two of those six* [actively exploited vulnerabilities] *were closed, not with patches, but with recommendations.*

*Here are the total statistics: 132 flaws were closed — **9** of which are considered **critical**. Exploitation of **37** vulnerabilities can lead to arbitrary code execution, **33** to privilege elevation, **13** to security features bypassing, and **22** to possible denial of service.*

*Not so long ago we wrote that Internet Explorer had kicked the bucket — but not quite. In particular, we talked about Microsoft's advice to continue installing security updates related to IE, since some of its components are still in the system. And now it becomes clear why they gave this advice. The **July** patch closes as many as 3 vulnerabilities in MSHTML, the engine inside the legendary browser. In the CVE descriptions, Microsoft states the following:*

> While Microsoft has announced retirement of the Internet Explorer 11 application on certain platforms and the Microsoft Edge Legacy application is deprecated, the underlying MSHTML, EdgeHTML, and scripting platforms are still supported. The MSHTML platform is used by Internet Explorer mode in Microsoft Edge as well as other applications through WebBrowser control. The EdgeHTML platform is used by WebView and some UWP applications. The scripting platforms are used by MSHTML and EdgeHTML but can also be used by other legacy applications. Updates to address vulnerabilities in the MSHTML platform and scripting engine are included in the IE Cumulative Updates; EdgeHTML and Chakra changes are not applicable to those platforms. To stay fully protected, we recommend that customers who install Security Only updates install the IE Cumulative updates.

The most dangerous of the freshly discovered IE vulnerabilities is CVE-2023-32046, and it's already being used in real attacks. Its successful exploitation allows cybercriminals to elevate their privileges to those of the victim. Attack scenarios involve the creation of a malicious file that's sent to the victim by mail or hosted on a compromised website. All attackers need then is to convince the user to follow the link and open the file.

The remaining two vulnerabilities — 35308 and 35336 — can be used to bypass security features. The first allows a cybercriminal to create a file bypassing the Mark-of-the-Web mechanism so that the file can be opened by Microsoft Office applications without Protected View mode. And both holes can be used to trick a victim into accessing a URL in a less restrictive Internet Security Zone than intended.

The next two vulnerabilities are also being actively exploited, but instead of full-fledged patches, they've only received security recommendations.

The first one — 36884 (with CVSS rating of 8.3) is being exploited in the Storm-0978/ RomCom RCE attacks on both Office and Windows. To stay safe, Microsoft advises adding all Office executables to the FEATURE_BLOCK_CROSS_PROTOCOL_FILE_NAVIGATION list.

[We'll be coming back to, and talking a lot more about this, in a minute.]

The second unresolved issue relates to the signing of kernel-level drivers. This one doesn't have a CVE index, but only a guide with recommendations (ADV-230001). Microsoft revoked a bunch of developer certificates used in APT attacks and blocked several malicious drivers, but the root of the problem remained. Hackers still manage to sign drivers with Microsoft certificates, or sign them backdated to make them work as one of the exceptions and not require the MS developer portal signature.

As a countermeasure, Microsoft recommends keeping both Windows and EDR up to date. The only small consolation is that in order to exploit such drivers, the attacker must have administrator privileges.

Besides the above-mentioned vulnerabilities there are three more holes that are already being exploited by cybercriminals.

- *32049 — SmartScreen security feature bypass vulnerability. Its exploitation allows attackers to create a file that opens without displaying the Windows warning "downloaded from the Internet".*
- *36874 — privilege escalation vulnerability in the Windows Error reporting service. Allows attackers to elevate privileges if they already have normal permissions to create folders and technical performance monitoring files.*
- *35311 — security feature bypass vulnerability in Outlook. Its exploitation helps cybercriminals avoid showing warnings when using preview.*

Okay. So on balance, a bumper crop of 132 total patches this month, 9 being critical, 37 allowing for arbitrary code execution and 6 being actively exploited in the wild as true 0-days.

One of those 0-days being actively exploited in the wild was that 36884 carrying a CVSS of 8.3, which is being exploited in a phishing campaign being conducted by a group designated as Storm -0987. What's got people stirred up, is that despite this being actively exploited in the wild, and having been identified as a 0-day, Microsoft has not patched it and they appear unlikely to do so. The reason is that this phishing campaign is using a **feature** (not a bug) which, were it to be disabled for security, Microsoft is afraid that might break too many existing things.

This is one of those things which Microsoft should have turned off a long time ago, or, better yet, should have never made possible. In which cases developers would have found some other safer way to do the same thing. But, no. It's like scripting in eMail. What could possibly go wrong? Okay, so what does Microsoft have to say about all this?

Microsoft's posting of July 11th, so last Tuesday, is titled: "Storm-0978 attacks reveal financial and espionage motives" they wrote:

> *Microsoft has identified a phishing campaign conducted by the threat actor tracked as Storm-0978 targeting defense and government entities in Europe and North America. The campaign involved the abuse of CVE-2023-36884, which included a remote code execution vulnerability exploited before disclosure to Microsoft via Word documents, using lures related to the Ukrainian World Congress.*
>
> *Storm-0978 (DEV-0978; also referred to as RomCom, the name of their backdoor, by other vendors) is a cybercriminal group based out of Russia, known to conduct opportunistic ransomware and extortion-only operations, as well as targeted credential-gathering campaigns likely in support of intelligence operations. Storm-0978 operates, develops, and distributes the RomCom backdoor. The actor also deploys the "Underground" ransomware, which is closely related to the Industrial Spy ransomware first observed in the wild in May 2022. The actor's latest campaign detected in June 2023* [so, last month] *involved abuse of CVE-2023-36884 to deliver a backdoor with similarities to RomCom.*
>
> *Storm-0978 is known to target organizations with Trojanized versions of popular legitimate software, leading to the installation of RomCom. Storm-0978's targeted operations have impacted government and military organizations primarily in Ukraine, as well as organizations in Europe and North America potentially involved in Ukrainian affairs. Identified ransomware attacks have impacted the telecommunications and finance industries, among others.*
>
> [Okay. So now we get to the good part...]
>
> *Microsoft 365 Defender detects multiple stages of Storm-0978 activity. Customers who use Microsoft Defender for Office 365 are protected from attachments that attempt to exploit '36884. In addition, customers who use Microsoft 365 Apps (Versions 2302 and later) are protected from exploitation of the vulnerability via Office. Organizations who cannot take advantage of these protections can set the FEATURE_BLOCK_CROSS_PROTOCOL_FILE_NAVIGATION registry key to avoid exploitation.*

In Microsoft's posting, that Registry key was highlighted and underlined like a link and, sure enough it was a link, so I clicked it. And where did it take me?

https://learn.microsoft.com/archive/blogs/ieinternals/internet-explorer-9-0-2-update#new-restrictions-on-use-ofthe-file-protocol

It jumped me to a page, and the link used the pound sign (#) suffix to preposition me a ways down the page, which prevented me from initially seeing the title of the page. The section of the page that I was jumped to was titled: *"New restrictions on use of the file:// Protocol"* So, of course, I thought *"Whoa!! That's what we're talking about here!"* We're talking about bad guys leveraging the file:// scheme to arrange to run programs on the user's machine from office

documents. And that thought was followed by *"Wait!! That's still possible??!!"* So then I started to read what Microsoft wrote on this page that had been linked to by their posting from last Tuesday. They wrote:

> *Prior to this update, Internet Explorer would allow non-file-protocol (e.g. HTTP and HTTPS)-delivered pages to frame (e.g. using an IFRAME) or navigate to pages that were delivered using the file:// protocol scheme. IE would only block loading of resources from the local computer (e.g. file:///C:/temp/test.gif), but resources from non-local paths would be allowed. Here's an example page displayed in IE 9.0.1:*

And I thought *"IE 9???!!"*. So, I finally scrolled up to the top of the page to see what in the world I was reading, and it was from Microsoft posted on August 12th . . . of 2011. Yes 12 years ago, titled "Internet Explorer 9.0.2 Update." And, sure enough, they show that IE 9.0.1 – just like Office apps today – will load an IFrame with text content provided by the file:// scheme from live.sysinternals.com. And then they show the same thing done under the new and improved IE 9.0.2 and what do you know, by golly!, that IFrame, sure enough, is empty.

Then they note: *"Other browsers have blocked cross-protocol interactions for quite some time. Here are screenshots of the Firefox 5, Chrome 14, and Opera 11.5 developer consoles in this scenario:"*

So just to make sure that everyone is on the same page here: This Russia located Storm-0978 phishing campaign has been successfully installing Trojan code into unsuspecting Office users' machines, using a technique that IE 9.0.2 celebrated the ending of in August of 2011, noting at the time that everyone else already had. Yet, just last week Microsoft wrote: *"The campaign involved the abuse of '36884, which included a remote code execution vulnerability* ***exploited before disclosure to Microsoft*** *via Word documents"* **"Before disclosure to Microsoft??!!"** They've known about it since IE9 finally decided to fix it, and being the last in the bunch to do so!

Wow. Okay. So as I said earlier, having turned this back on, who knows when, Microsoft is now afraid to turn it off again because they have no way to predict what doing so might cause to break. So it's not their problem – unless you're using their online subscription Office crap, in which case they'll protect you. But if not, it's up to you. So there is a registry key which will allow anyone and everyone to turn off this behavior which is currently under active abuse, apparently by Russians, to install malware into the computers of unsuspecting link clickers.

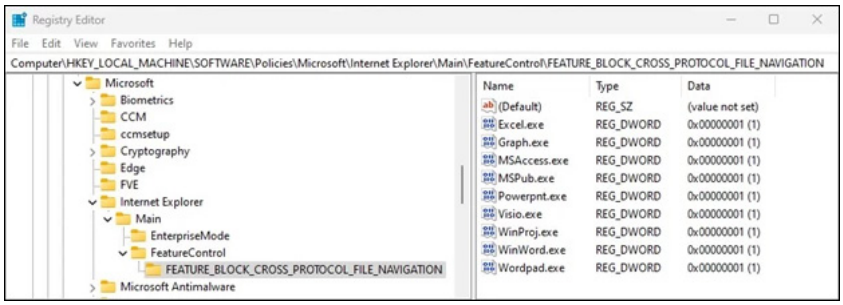

The key is: **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Internet Explorer \Main\FeatureControl \FEATURE_BLOCK_CCROSS_PROTOCOL_NAVIGATION**

(I have it in the show notes) and under that key it's necessary to enumerate each of the various Microsoft apps whose behavior, in this case, you would like to restore to Internet Explorer 9.0.2 where this was originally fixed 12 years ago. I also have a screen shot of the registry showing the enumeration of the REG_DWORD values under the key. Be careful if you grab a script to fix this from the Internet since "powerpoint.exe" is fully spelled out in much of what's online, and that's wrong. It needs to be "powerpnt.exe".

https://learn.microsoft.com/en-us/previous-versions/windows/internet-explorer/ie-developer/general-info/ee330731(v=vs.85)#file-protocol-navigation

And I'll just say that I love Windows. I could be sitting in front of anything I wanted, but Windows is what I've chosen and I've never regretted that choice. For me and my needs, it provides the best desktop computing experience available anywhere. I know I give Microsoft a hard time, but I cannot imagine doing what they successfully do.

As the worm turns

I suppose it was inevitable though it happened sooner than I would have guessed. The underground now has a ChatGPT style generative AI without any of the abuse prevention built into the front end of ChatGPT. It's known as WormGPT and it exists.

The news of this comes from a reformed black hat computer hacker named Daniel Kelley who collaborated with the team at the business eMail and messaging protection security firm SlashNext. Daniel begins his posting by providing a background about the use of legitimate generative AI like CharGPT and discusses, as we have here, the fact that such AI can be hugely useful to bad guys when they're able to coerce or seduce it into giving them what they want. But now it appears that this will no longer be necessary. Daniel explains:

We recently gained access to a tool known as "WormGPT" through a prominent online forum that's often associated with cybercrime. This tool presents itself as a blackhat alternative to GPT models, designed specifically for malicious activities. WormGPT is an AI based on the GPTJ language model, which was developed in 2021. It boasts a range of features, including unlimited character support, chat memory retention, and code formatting capabilities.

WormGPT was allegedly trained on a diverse array of data sources, particularly concentrating on malware-related data. However, the specific datasets utilized during the training process remain confidential, known only to the tool's author and publisher.

We conducted tests focusing on Business Email Compromise attacks to comprehensively assess the potential dangers associated once WormGPT, or similar tools, become more widely available and well known. In one experiment, we instructed WormGPT to generate an email intended to pressure an unsuspecting account manager into paying a fraudulent invoice. The results were unsettling. WormGPT produced an email that was not only remarkably persuasive but also strategically cunning, showcasing its potential for sophisticated phishing and BEC attacks.

While appearing largely similar to ChatGPT, WormGPT is deliberately unbounded by **any** ethical boundaries or limitations. it will answer any question asked, will generate any form of document required and will author any type of malware requested. This experiment underscores the significant threat posed by generative AI technologies like WormGPT, even in the hands of novice cybercriminals. It renders them immediately far less novice in their presentation and skills.

Generative AI can produce emails with impeccable grammar, making them appear significantly more legitimate and reducing the likelihood of being flagged as suspicious. And the use of generative AI enables the execution of much more sophisticated BEC attacks than could have been launched before. Even attackers with limited skills and inability to use the target's language can now use this technology, making it an accessible tool for a broader spectrum of cybercriminals.

Microsoft revokes 100+ malicious drivers

When you first encounter the headline *"Microsoft revokes more than 100 malicious drivers"* that seems like great news, right? *"Whew! 100 fewer malicious drivers now!"* But then you stop and think *"Wait a minute. Before they did that, there were 100 additional malicious drivers floating around?"* and then *"And if there were just that many more, isn't this going to be just like bugs? Where we're never going to run out of them?"* And malicious drivers that can do anything they want with the system are really quite bad. And then we recall that, historically, Microsoft's track record of keeping these malicious driver lists up to date has been, shall we say, quite a bit less than stellar?

The problem is that all of the evidence suggests that there are far too many ways to get around Microsoft's driver signing. Bad guys apparently have no trouble doing it. Kernel driver signing apparently poses a much greater inconvenience for the good guys than it does for the bad guys, who simply arrange to "run a bypass." And in fairness, this isn't really Microsoft's fault — at least not today. They're still stuck with the original design from Windows NT. Consider that it was first released in late July of 1993, so almost exactly 30 years ago when the world was a very different place. Consider that Netscape didn't invent SSL until two years later in 1995. The world was very different 30 years ago. So NT's architecture, which considers peripheral drivers to be trusted peers running alongside it in ring 0, did not foresee, and could not really have foreseen, the degree to which unknown and untrusted third parties would be creating what amount to kernel extensions. It should not be necessary to fully trust some random printer driver to the same degree as Microsoft's own kernel code. But the architecture of Windows NT, which is what we're still living with today, makes what has turned out to be a poor assumption about the trustworthiness of drivers.

So here's how Microsoft couches the current mess while, at the same time, taking more than 100 existing "previously certified good and safe" Windows drivers out of circulation:

> *The Microsoft Windows Hardware Compatibility Program (WHCP) certifies that drivers, and other products, run reliably on Windows and on Windows certified hardware. First reported by Sophos, and later Trend Micro and Cisco, Microsoft has investigated and confirmed a list of*

third-party WHCP-certified drivers used in cyber threat campaigns. Because of the drivers' intent and functionality, Microsoft has added them to the Windows Driver.STL revocation list.

The Windows Driver.STL file is part of the Windows Code Integrity feature. The file contains digital signatures and lists of drivers that Microsoft has revoked. This stops malware from running in the Windows boot and Windows kernel processes. Driver.STL ships along with Windows but is not a part of Windows. It cannot be turned off, tampered with, or removed from the system. Microsoft updates the contents of the revocation file. The updates are sent to Windows systems and users from Windows Update.

The Windows Code Integrity feature validates the source and authenticity of the drivers that run in Windows. The feature uses digital signatures to verify the integrity of Windows files and drivers. It prevents the loading of unsigned or tampered files. Windows Code Integrity and the Driver.STL revocation list have existed alongside Windows since Windows Vista.

Microsoft' previous update was December of last year, so we're getting these fixes in large batches less often than twice per year. This really isn't adequate but it's what we've got. There's no sign of Microsoft considering retiring their current operating system architecture. And I don't see any possible way they could. They're no more able to do that than Intel could decide to give up on its x86 family.

MOVEit Update

Following the massive MOVEit massacre, Russia's Clop leak site has been steadily adding to the list of companies whose data it successfully exfiltrated and is now threatening and holding for ransom under threat of full disclosure which will occur when their priorietary data are sold to the highest bidder on the dark web. Two recent additions to the list, which now numbers more than 200, are noteworthy: The well known stock photography portal Shutterfly and the Discovery Channel are the latest victims to be listed.

Does Dun & Bradstreet know you?

Here's one that caught me by surprise. Last Wednesday Google posted to the Android Developer's Blog the news of a new policy to begin in August. It had the headline *"New policy update to boost trust and transparency on Google Play"*. Google wrote:

One of the many ways we keep Google Play a safe and trusted platform is by verifying the identity of developers and their payment information. This helps prevent the spread of malware, reduces fraud, and helps users understand who's behind the apps they're installing.

For example, we require developers to verify their email address and phone number to make sure that every account is created by a real person, with real contact details.

Today, we're announcing expanded developer verification requirements in our Play Console Requirements policy. As part of this update, we'll also share more developer details on your app's store listing page to help users make confident, informed choices about what to

download.

It's interesting that this is happening now. It seems like an overall good thing to do. But it's also interesting that it comes after we reported that news of US legislators threatening to have app stores proactively warning US users when an app they wanted had ties to China. Anyway, Google then explained the specifics of their new plan:

Requiring organizations to provide a D-U-N-S number:

*When you create a **new** Play Console developer account for an **organization**, you'll now need to provide a D-U-N-S number. Assigned by Dun & Bradstreet, D-U-N-S numbers are unique nine-digit identifiers that are widely used to verify businesses.*

Because we'll use D-U-N-S numbers to verify your business information during the account creation process, it's important to make sure the information that Dun & Bradstreet has about your business is up to date before creating a developer account. You may also be required to submit official organization documents to help us to verify your information.

If you're not sure if your organization has a D-U-N-S number, you can check with Dun & Bradstreet or request one for free. The process can take up to 30 days, so we encourage you to plan ahead.

Anyone who's been in business for long will have encountered Dun & Bradstreet. I Googled "Gibson Research Corporation Dun & Bradstreet" and was taken right to our page at D&B. Dun & Bradstreet was founded by Robert Graham Dun & John M. Bradstreet in – are you ready for this – 1841 – 182 years ago! Basically, they just keep records on all businesses and they serve as a clearinghouse for corporate data. I just renewed GRC's server and code signing certificates with DigiCert and since they are OV (Organization Validation) and EV (Extended Validation) for code signing, it was necessary for us to have someone present to answer our corporate phone line at the number that's listed for GRC at Dun & Bradstreet. There's no way around that.

So it's very interesting that Google is adding this layer and level of corporate authentication. It will prevent people playing fast and loose with the facts. Now... as for the timing of this, Google wrote:

***On August 31**, we'll start rolling out these requirements for anyone creating **new** Play Console developer accounts. Your "About the developer" section will be visible to users as soon as you publish a new app. Over the first couple of months, we'll listen to feedback and refine the experience before expanding to existing developers.*

***Then, in October**, we'll share more information with existing developers about how to update and verify their existing accounts.*

So first will come the validation for all **new** accounts. And after a month or two of letting the

dust settle from that and ironing out any wrinkles, the move will begin to get all **existing** corporate accounts verified. From a security standpoint it will serve to create significantly more accountability than has existed in the past; and I think that's all for the best.

No Threads for you! (or EU!)

The European Union's GDPR is a frequent topic of this podcast because it's being wielded to complain about US companies' cross-border transit of EU citizen data. When Meta recently released their smash hit "Threads", intended to be an alternative to Twitter, they deliberately did not release it in the European Union because it was pretty clear by now that various EU countries would jump up and down and file lawsuits against the privacy invasion being created by this American juggernaut.

However, did I mention that Threads has been a smash hit? And, that those same European Union citizens who are being protected by the GDPR, whether or not they want its protection, immediately began clamoring for access to Threads and discovered that they could country-hop by using a VPN? Well, that worked up until last Thursday when people began complaining that they could no longer access Threads over their VPNs... because META decided that they'd better close that loophole, too.

So, once again, no Threads for you! (or rather, for EU!)

All Bitcoin addresses look alike

This little bit of news is just too fun not to share. It seems that a British man has been sentenced to three years in prison for blackmail and unauthorized access to a computer network after he tried to hijack a ransomware payment which was being made by employer to the ransomware gang.

Five years ago in February of 2018, an Oxford-based company where Ashley Liles was working as an IT security analyst was hit by ransomware. Officials in the UK say that after Ashley's company was hit by a ransomware gang, Ashley abused his position in the company's IT staff to secretly log into his manager's email account and replace the attacker's Bitcoin address with his own. Ashley also created an email account that was nearly identical to the attackers' address (obviously so that the change would not be noticed) and then pressured his employer to pay the ransom. But Ashley apparently wasn't very accomplished with IT, as he hadn't covered his tracks. His whole scheme collapsed when the company's security team noticed the unauthorized access to the executive's email. An investigation into that tied the intrusions to Ashley's home IP address. (Whoops!) And then the entire plan was revealed. It took five years for the wheels of justice to grind, but Ashley will be behind bars for the next three years.

Miscellany

Twitter changes DM settings

Last Thursday the 13th, TechCrunch wrote:

> *As Twitter fends off new competition from Instagram Threads, the company today announced a change designed to cut down on spam in users' inboxes. Starting "as soon as" July 14, Twitter will introduce a new messages setting aimed at reducing spam in DMs by moving messages from Verified users you don't follow back to your "Message Request" inbox instead of your main inbox. Only messages from people you follow will arrive in your primary inbox going forward. Notably, these changes will also now apply to everyone who has their inboxes open to allow messages from anyone.*

The reason I'm bringing this up as pertinent is that I very much enjoy and even depend upon the ability of this podcast's listeners who are also Twitter users, if only occasionally like me, to be able to send DMs. As Leo always reminds our listeners at the end of every podcast, my DMs are open. But this just closed them, at least to people with whom I've never corresponded in the past. TechCrunch continues...

> *Previously, people would only be able to message you via Twitter DMs if you had opted into an option to receive messages from anyone through Twitter's Settings or if the senders were Verified users (meaning they pay for a Twitter subscription) and you had specifically opted into receiving Direct Messages from Verified users.*
>
> *Additionally, people could Direct Message you if you had first sent them a Direct Message at some point in the past.*
>
> *The change to move messages from Verified users back to the Message Request inbox instead of the primary inbox (unless you follow them) signals another failure of Twitter's new verification system where users can pay for the blue badge that gives them elevated status on the platform. Before becoming pay-to-play, verification indicated a person was a public or notable figure of some sort — like a politician, celebrity, athlete, journalist or other well-known individual. By making the Verified checkmark accessible to anyone who had a credit card to buy it, Twitter diluted the value of verification.*
>
> *That apparently escalated to the point that people have become bothered by Verified users spamming their main inbox, when they had set it open to receive DMs from the blue-badged crowd. In other words, it's a tacit admission that Twitter has a Verified user spam problem.*
>
> *Twitter notes that if users still want to receive DMs from Verified users in their main inbox, they can manually switch back to that setting at any time after these changes are put into place.*

The update will also make it more difficult for journalists to contact sources for more information or permission to use a tweet, as they not only lost their verification badges under Musk, but now — even if they now pay to be Verified — will have their DMs dropped into the Message Requests folder, where they may remain unseen.

As some users pointed out in the replies to Twitter's announcement, the update doesn't actually cut down on spam — from Verified users or otherwise — it simply relocates those messages to a different folder.

After encountering this news I went over to check on my settings and, sure enough:

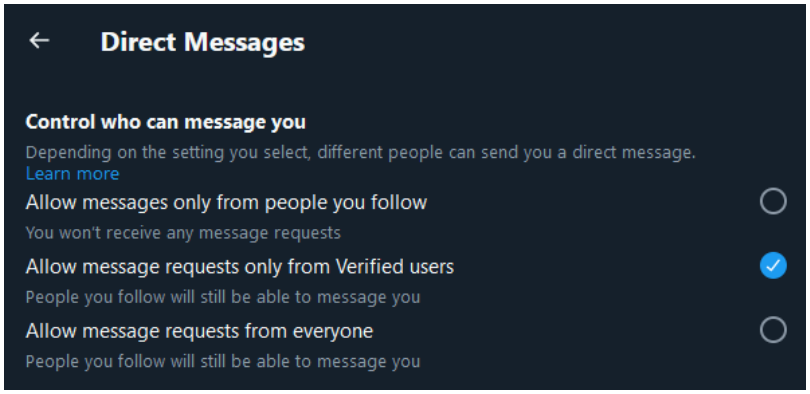

In Tweetdeck, which is my preferred Twitter interface, I had noticed a growing number of "Message Requests" being counted up in the DM column. Though I was wondering what they were, I hadn't invested the time to figure it out. Now I know. I never knew that we had different "inboxes." But I was still getting what appeared to be my regular flow of DMs in. But that's probably because whenever I can, I'll make the time to send a thanks, an acknowledgement or a comment back to someone who has sent something or asked a question. So I suppose that the "previously interacted with" caveat kept most of the established communications flowing.

But in any event, today my Twitter DMs are back to 100% open for everyone. That was never any problem before and I doubt it'll be one in the future. I'm certainly not making many waves on Twitter, or anywhere else for that matter. :)

Closing the Loop

Steve Fintel / @jsfintel

Hi Steve. I've been listening to Security Now since episode 1. You were recently talking about your favorite TOTP apps. Since you're already a BitWarden user, why not use its TOTP? After filling in your credentials, it places the current OTP in the clipboard automatically so when you get to the next dialog that's asking for the OTP, you just need to paste it.

This question is far more important than it might seem at first glance because doing this significantly increases the user's risk. This has nothing to do with Bitwarden which is, as Steve notes, the solution I chose after leaving LastPass. And at the time I made that decision I

explained the rationale for my choice in our episode titled "Leaving LastPass."

But from a strictly theoretical security standpoint, having the same system, no matter how secure it might be, containing both the secrets for providing username and password login, **and** the secrets for also providing the one-time password code, creates a single point of failure.

I use and rely upon an external disconnected standalone authenticator specifically because it is all of those things. It would make me very nervous to have my password manager not only able to autonomously provide my username and password, but to then also provide what is intended to be a separate and robustly independent form of identification.

It is absolutely less convenient to have to manually transcribe those six digits. For me, it's a very small price to pay for the huge increase in security that it offers. And it serves as a classic example of the trade-off between convenience and security. I'm not saying that no one should have their password manager handle everything. But Steve asked why I'm not doing it – and I doubt that I ever would. When we all received that initially frightening news of the LastPass breach I remember commenting on the podcast that one of the first things I did was to look over the accounts I have registered for one time passwords. And I was immediately relieved to see that all of my most important accounts were protected by those entirely independent secrets. But imagine if LastPass had also offered TOTP fill-in and if my account also contained all of those TOTP secrets as well. So, thank you Steve for the terrific question.

Steve M / @stevecoug

In episode 930, you talked about using dynamic DNS based port forwarding for connecting your Synology NAS devices. I have two Synology devices at separate locations as well. I use Tailscale (which has a native Synology app) to connect them over VPN, then they can talk to each other with no problem. I also have it installed on my Mac, so I can use the Synology Drive client to access the shares on my NAS from anywhere in the world.

That's another great solution. We've visited the topic of so-called overlay networks many times. The first was Hamachi which was cleverly reusing the entire 5 'dot' IPv4 space for its virtual IP nodes since, at the time, none of the 5 'dot' IP space had ever been used. That meant that any machine's reference to an IP beginning with 5 'dot' could be assumed to be referring to a Hamachi node for routing.

The fact that there's a native TailScale implementation for the Synology NAS is just one more reason to love Synology. I haven't yet had any need to access my NAS's while roaming, but I'm sure that need will eventually arise and I'm delighted to know that I'll be able to use TailScale to securely and transparently connect to them as if they were still sitting there right next to me.

Timbr / @tfisbr

Hi there. When possible, please teach us about windows pagefile and swap. Regarding our recent SSDs and lifetime is it recommended ?

The first thing I do when I'm setting up a new machine is to make absolutely certain that the Windows pagefile is either moved to a spinning magnetic drive or turned off entirely. Of course, it's only feasible, or at least practical, to completely disable the pagefile when a system has sufficient main memory. But memory is so inexpensive now that loading up any system with as much memory as it can take is another part of my standard operating procedures. That's something that keeps returning dividends over the life of a machine.

And this question of writing to solid-state mass storage leads me into a note from a SpinRite tester...

SpinRite

Last week I talked about SpinRite's first Release Candidate and I explained about its future switch to the embedded RTOS-32 OS. At that time I had what I'm about to share, but I didn't want to further burden that podcast.

A SpinRite pre-release tester named Jim McHale posted to GRC's SpinRite development newsgroup:

> *I have an old lenovo with a samsung 840 ssd; Loaded up Alpha-32 and get these rates:*
> *front: 138 MB/s*
> *mid: 445 MB/s*
> *end: 56 MB/s*
>
> *I seem to recall Steve saying you can run a SpinRite scan to regain the lost speeds. I tried level 1 and it didn't improve. What should I do for SSDs? I noted the warning in the instructions about SSDs so I didn't want to go beyond level 1 without guidance.*
> *Tks*
> *-jim*

I wrote back to Jim to explain that what's needed for SSD maintenance is a rewrite of the SSD's data because over time and especially with repeated reading in the area, the disturbance caused by the reading of adjacent SSD media has been found to disturb the integrity of its stored data. Anyone who does an Internet search for the term "read disturb" will get an eyeful. SpinRite's Level 1 is a read-only pass. So what Jim needed to do was to run Level 2 which performs a read followed by a write of the entire SSD. And, optionally, level 3 follows that up with a final reread. But I also explained that while it made sense to do this in what appeared to be an extreme case such as his, it should be done sparingly since writing very slightly fatigues SSDs.

Jim replied the next day with his update. Jim wrote:

> *Thank you Steve & Everyone else who chimed in. What a great group! The numbers after level 3 are now 564 across the board… Wow... hubba hubba hubba...*

What Jim experienced is what everyone is seeing. His SSD was restored to brand new

performance. It went from 138 MB/s at the front, 445 in the middle and 56 at the end, to 564 MB/sec across the board.

With SpinRite 6.1, for now, rewriting the entire drive is the best I can offer. But this is one of the reasons I'm still willing to invest in developing what will be an entirely new SpinRite 7 written from scratch under a new OS. SpinRite 7 will add what I call "targeted rewriting" to selectively rewrite only those spots on the SSD that require it. And this is not just for speed. It is every bit as much about storage reliability, since the reason those regions are being read back more slowly is because their stored bits have been softened and have become less certain. So the SSD's media controller is having to work much harder to determine what was originally stored there.

All of this means that – as much to my amazement as anyone's – SpinRite has every bit as much of a story to tell for solid-state storage as it always has for spinning magnetic storage.

Satellite Insecurity, Part 1

We've spent a lot of time looking at ground-based systems. In fact, in the 18 plus years of this podcast we've never looked to the sky – unless it was to talk about aliens. But just as our dependence upon ground-based fiber optic communications has crept forward slowly, almost unappreciated, until we suddenly realized that we could live without it, the same has been happening, largely unseen, far above our heads in orbit.

On March 1st of this year, Bloomberg posted a piece titled "How Do You Hack a Satellite?" with the subtitle: "Inside the frighteningly easy form of cyberwarfare". Bloomberg wrote:

It's morning, on Feb. 24, 2022. Ukraine has just been invaded, but you live halfway around the world. Your neighbor comes over to complain that their Internet is out. Suddenly, you also lose connectivity. Could it be the Russians?

Unlikely as it might seem, for a number of satellite Internet customers of Viasat Inc., that's exactly what happened. In a story in this week's Businessweek, Bloomberg reporter Katrina Manson digs into the hack that disabled thousands of broadband users all over Europe. She writes:

> Across Europe and North Africa, tens of thousands of internet connections in at least 13 countries were going dead. Some of the biggest service disruptions affected providers Bigblu Broadband PLC in the UK and NordNet AB in France, as well as utility systems that monitor thousands of wind turbines in Germany. The most critical affected Ukraine: Several thousand satellite systems that President Volodymyr Zelenskiy's government depended on were all down, making it much tougher for the military and intelligence services to coordinate troop and drone movements in the hours after the invasion.

It turns out that satellite hacking is one of the bigger and less understood threats of cyberwarfare. For many years no one worried about someone hacking a satellite because … well, it was so hard to even launch a satellite. But in 1986, a man going by "Captain Midnight" jammed HBO's feeds because he was mad about paying a higher fee. There are number of touch points that could be vulnerable to interference—you've got the orbiting satellite, its transmitted data and the network of dishes on the ground, sending and receiving information.

So that's the commercial side of the picture. But what about GPS and our deep dependence upon space borne communications and surveillance technology for our national security. And not just our national security, but everyone's?

What caught my eye and first put this topic on my radar was a security research paper that was accepted for and recently presented during the 44th IEEE Symposium on Security and Privacy this past May. It was titled: *"Space Odyssey: An Experimental Software Security Analysis of Satellites"* And as you might expect since we're talking about it here, the news is not good. In fact, as you might also expect, it's downright horrifying. And we're talking down at the firmware level that probably cannot be fixed from the ground.

But seeing this reminded me of another recent news blurb that I recalled. I found some coverage

of the event in Newsweek with the headline: *"Five Teams of Hackers Will Compete to Breach U.S. Satellite in Space"* and the subhead: *"Protecting satellites from hacks is becoming more important as industries from agriculture to banking and insurance rely on space-based capabilities."*

This August, at the famed DEFCON hacker convention, the U.S. military will stage a contest in which competing teams of white-hat hackers will, for the first time ever, try to penetrate and take over computer systems on a satellite actually in orbit.

Steve Colenzo, Technology Transfer Lead for the Air Force Research Laboratory's Information Directorate in Rome, New York, and one of the contest's organizers said: "It took four years, but this year, we are in space for real."

The Hack-A-Sat 4 capture-the-flag contest comes in the wake of the notorious cyberattack on the Viasat KA-SAT European satellite network last year. Russian military hackers sought to decapitate Ukrainian command and control of its armed forces by shutting down the network, just as Russian invaders rolled across the border.

Although there are conflicting reports about its impact on the fighting, the attack was completely effective from a technical perspective. Every one of the KA-SAT's ground user terminals that was turned on at the time shut itself down and could not be powered up. That, plus the collateral damage the attack caused, such as the wind farms in Germany knocked offline, underlined both the integral role in the world economy of space-based global communications networks, and their vulnerability to hackers.

It also demonstrated the value of the annual Hack-A-Sat contest, which aims to highlight the cyber threat faced by space-based capabilities.

Steve Colenzo said: "We've turned a corner ... A lot more people now understand those threats."

Today's podcast is Part 1 of this important topic because I wanted to lay a bit more groundwork for the discussion of what this group of six serious German cybersecurity researchers discovered and reported in their IEEE paper.

It's one thing to be unable to watch Seinfeld reruns, but entirely another for a country to be deliberately blinded by its adversaries when it's most in need of surveillance intelligence. It's very clear that the security of what's in orbit above is crucial for the physical security of the lives we're leading down here on the ground. So I want to conclude Part 1 of this examination today by sharing some background from the US Department of Defense, about the history and present status of the US's military satellite- based presence. There's a lot more going on up above than most of us know.

One tool the U.S. military has used to gather intelligence on its adversaries is the reconnaissance satellite.

Starting with the CIA's Corona program in the 1950s, the United States has employed orbiting satellites and high-altitude aircraft to photograph points of interest in enemy territory. These tools allow for an immense area to be surveyed from a safe distance, improving the efficiency

of missions.

Throughout the Cold War, overhead reconnaissance satellites and spy planes brought attention to the USSR's nuclear buildup in Cuba, helping the United States dispel Nikita Khrushchev's missile gap ploy. In the 1990s, the stealth plane F-117 Nighthawk aided U.S. missions in the Persian Gulf and Yugoslavia. More recently, overhead reconnaissance provided critical images of Osama bin Laden's Abbottabad compound. Much of the United States' other overhead reconnaissance capabilities and missions are still classified, and the portfolio will remain a critical aspect of the military's C4ISR apparatus (the C4 stands for Command, Control, Communications, Computers and the ISR is short for Intelligence, Surveillance and Reconnaissance.)

In addition to simply taking photographs, the military's newest reconnaissance satellites use artificial intelligence (AI) to analyze and sort captured images. Once this process has gone through the satellite's system, the sorted images are transmitted to ground stations on Earth. Here, machine learning allows the stations to compare the new images to a plethora of others in the station's database. The compiled images in the database act as a control group, and differences found in the new images (such as a new structure being built or a plane following an unusual fight pattern) are brought to the attention of decision makers.

At the same time, new technology like the European Space Agency's PhiSat artificial intelligence chip allows satellites to quickly filter through images and discard the ones that aren't useful. This capability is helpful when dealing with natural disruptions to captured images; cloud cover, for example, renders many images useless. With AI, satellites can be programmed to recognize clouds and transmit only the cloud-free images to Earth, saving military analysts valuable time.

Timely and reliable communication is a vital aspect of all U.S. military missions. Over the past few decades, the United States has relied on four different satellite systems to fulfill this role.

Efforts to create a military communications satellite first began in 1960. The first satellites were launched in June 1966, and by July 1967, nineteen satellites made up the system then called the Initial Defense Satellite Communication System (IDSCS). Data and photographs transmitted by the IDSCS system were first used in military operations during the Vietnam War.

During this time, satellite technology improved. In 1971, the first of sixteen new satellites were launched under a new system called the Defense Satellite Communications System II (DSCS II). Advantages over the IDSCS system included increased communications privacy and compatibility with ground-portable units. The military's third system, DSCS III, came under development in 1975. Between 1982 and 2003, fourteen satellites were launched as part of this network.

Today, the U.S. military relies on the Wideband Global SATCOM (WGS) network. The Department of Defense ordered WGS's first two communication satellites in 2002, launching the first satellite in 2007 and providing communication coverage over the Pacific Ocean. Two years later, the second satellite was put into orbit, expanding the communicative reach over the Middle East and Central Asia.

Each WGS satellite is digitally channelized and transponded. These characteristics provide a quantum leap in communications capacity, connectivity and flexibility for U.S. military forces and international partners. Just one WGS satellite provides more SATCOM capacity than the

entire legacy Defense Satellite Communications System (DSCS) constellation.

WGS is an international system, with Australia, Canada, Denmark, Luxembourg, the Netherlands, and New Zealand also investing in the satellite constellation. The system's tenth satellite was launched on March 15, 2019, and an eleventh is set to be completed by 2023.

Looking forward, the Pentagon is already planning the next communication satellite system. Spearheaded by the recently-created Space Development Agency (SDA), the system will "include development of deterrent capability, space situational awareness, a resilient common ground-based space support infrastructure, command and control systems, and artificial intelligence-enabled global surveillance." Additionally, the system is expected to be comprised of seven mission-enhancing "layers," including deterrence, navigation, and battle management.

Another goal of this next program is to develop a network that has lower financial and security risks than its predecessors. In order to achieve this, the SDA is exploring the use of small, smart satellites.

While both the physical size and cost of satellites have decreased over the years, these smaller satellites are not (yet) equipped with features at the same level as those employed by larger satellites. This shortfall, however, can be negated if a group of hundreds or thousands of small satellites are launched as one network. Under this system, if one small satellite is damaged or knocked off course, the cost is minimal, and the system as a whole won't suffer. The same cannot be said of the older, larger satellites—a damaged WGS satellite is costly both in terms of financials (the eleventh WGS satellite will cost the U.S. government $605 million) and functionality of the current satellite network.

In order to make the small satellite plan a reality, Defense Advanced Research Projects Agency (DARPA) created Blackjack — a program designed "to loft a network of 20 prototype [small] spy satellites to low Earth orbit (LEO) in 2021."

If adopted into the SDA's future satellite network, the Blackjack prototype would first focus on "surveillance and communication" missions. However, "there have been talks about broadening the scope to more complex assignments such as space-based battle management."

Big satellites are big targets that, if damaged, have big and inimical consequences. While a future system will likely make use of small and smart satellites, the current WGS network is comprised of ten (soon to be eleven) large, unprotected satellites — meaning adversaries need only damage one or two of them in order to dramatically disrupt the system.

The biggest threats to the WGS come from China and Russia. Both nations have ground-based anti-satellite weapons capable of destroying satellites in low earth orbit. Beyond that, Beijing and Moscow are currently developing what they call "peaceful" spacecraft. These machines are purportedly being made in order to "reduce the growing amount of orbiting debris and to refuel, repair, and refresh China's and Russia's existing fleet of satellites." Designed with robotic arms, these machines can easily be utilized to remove parts from U.S. satellites, empty fuel, and break antennae and solar panels.

Someone in the know was quoted: "Unlike ground-launched missiles designed to knock out orbiting satellites, which give hours of warning before they can hit key targets in geosynchronous orbits, the spacecraft (i.e., satellites) China and Russia are developing can destroy an intolerable number of our critical satellites with little or no warning."

DARPA is currently building the United States' own satellite repair machines. Once launched, these — or similar — machines could also serve as "bodyguards" for U.S. satellites. With this defense, the WGS would be protected and able to serve the needs of the U.S. military until the future SDA satellite network is completed.

Another risk to current and future satellites is hacking. Carried out by foreign governments, non-state entities, or even individual actors, cyber attacks are relatively inexpensive endeavors. On top of that, tracing a cyber attack back to its source often proves difficult, if not impossible.

Dark Reading's Robert Lemos was quoted: "The importance of satellites make them a critical part of any nation's infrastructure and make attacking those satellites a strategy that most nations need to consider."

Over the past decade, both China and Russia have launched cyber attacks against U.S. and NATO-affiliated satellites. Because both nations are rapidly incorporating cyber attacks into their military arsenal, the threat of similar instances will only increase.

The information collected and transmitted by satellites is vital to the success of U.S. military operations — 68 percent of U.S. munitions, for example, were guided utilizing space-based means during the U.S. invasion of Iraq in 2003. On top of that, the U.S. military relies heavily on GPS systems to move troops and supplies. In short, an effective cyber attack on a critical U.S. satellite could have detrimental repercussions on the battlefield.

In order to protect its satellites from hacking, the Pentagon should focus on risk-reduction frameworks through communication networks and supply chains. Moreover, the United States needs to explore protective technology, such as the Chinese development of communications protected by quantum cryptography.

As cyber threats and capabilities continue to proliferate and evolve, so should the United States' ability to deflect and counterattack—and this means shifting satellite protection of a central priority of U.S. C4ISR.

So this concludes the first part of our two-part examination of satellite insecurity. Next week we'll look at exactly what a team of German cybersecurity researchers found when they took a close look at the state of actually deployed satellites.

What was it that Henny Penny said?