The NSA @ Home

Description: What mistake did Windows Update make last week? What if you don't want to paste with formatting? What browser is building in a limited-bandwidth VPN? What more did we just learn about LastPass's second breach? What did Signal say to the U.K. about scanning its users' messages? What was just discovered hiding inside the Python package Index repository? What proactive move has QNAP finally taken? What disastrous bug did SpinRite's testers uncover last weekend in motherboard BIOSes? And what amazingly useful "Best Practices" advice has the NSA just published for home users? Answers to all those questions and some additional thoughts will be yours before you know it on this week's 912th episode of Security Now! titled "The NSA @ Home."

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-912.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-912-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. A mistake in Windows Update? Steve explains. We'll also talk about LastPass, now that more details are coming out. It's really kind of a stunning hack. Signal says ta-ta to the U.K. Well, it will if the U.K. does something bad. And believe it or not, the NSA's security recommendations. They're pretty good. All that and more coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 912, recorded Tuesday, February 28th, 2023: The NSA @ Home.

It's time for Security Now!, the show where we cover the latest in security with the most important guy in this building right now, Mr. Steve Gibson - and he's not even in the building.

Steve Gibson: Hey, Leo.

Leo: Hey, GRC.com.

Steve: Good to be with you again.

Leo: Yeah, it's good to see you.

Steve: Yeah.

Leo: Time once again to talk about how bad the world is.

Steve: Oh, we squeaked one more episode in on this last day of February. Of course, February ended early, so that's what made it squeaky tight. So we've got a bunch of questions in our new style that we're going to answer, and I really do wish that our audience had heard you encounter our Picture of the Week for the first time because...

Leo: Oh, it's a belly laugh, wow.

Steve: It is. Every time I look at it, it cracks me up again. But we're first going to answer some questions. What mistake did Windows Update make last week? What if you don't want to paste with formatting? What browser is building in a limited bandwidth VPN? What more did we just learn about LastPass's second breach? What did Signal have to say to the U.K. about scanning its users' messages? What was just discovered hiding inside the Python Package Index repository?

What proactive move has QNAP finally taken? What disastrous bug did SpinRite's testers uncover last weekend in motherboard BIOSes? What amazingly useful best practices advice has the NSA just published for home users? Answers to all those questions and some additional thoughts will be yours before you know it on this week's 912th episode, still going strong, of Security Now! titled "The NSA @ Home."

Leo: Ooh. Oh, that's going to be interesting.

Steve: Bring your NSA home with you.

Leo: Yeah. And that LastPass thing, holy cow, what a revelation.

Steve: Uh, yeah.

Leo: Finally they gave us some details, and it's not good.

Steve: Yup.

Leo: It's bad. It's very bad.

Steve: You know, I had the feeling, you don't want these bad guys on your tail because, I mean...

Leo: Yeah, no kidding. This was a very aggressive attack. They really knew what they wanted. Which kind of does not bode well for those of us whose LastPass vaults are now in their hands.

Steve: Yeah.

Leo: I mean, they went after them. It wasn't some script kiddie accidentally found them. Steve, talking about zero knowledge, our Picture of the Week.

Steve: What we have is a close-up photo that I presume someone took when they were checking out the smoke alarm on the home they just bought, or maybe the apartment they're renting, I don't know. So we've got a First Alert smoke alarm in close-up. Over on the left you see the little door you can pull off, probably to change the nine-volt battery and so forth.

Anyway, the focus of this picture, and the intent of taking a picture of the smoke alarm, is that printed on the side is a field where you can indicate the date at which this was installed because, you know, that could be important. Like fire extinguishers have an expiration date. You need to change them out every so often. So this says "Installed on:" and then there's an underline where clearly you're intended to indicate the date of installation. What we have instead written here is "Installed on: the ceiling."

Leo: Oh, boy.

Steve: Oh, that's where it is? That's why I'm on a ladder? Okay, yeah. Installed on the ceiling. I'm not sure where else you would install it. On the wall? On the floor? I don't...

Leo: I think, though, inspection will reveal that it is in fact installed on the ceiling. No need to write that down.

Steve: That's true, yeah.

Leo: You're looking at it.

Steve: You don't want to leave that field blank, however. Maybe it won't work without something written there. So, let's see, what could they possibly...

Leo: Installed on.

Steve: What could they want me to...

Leo: No, it's obvious. They want to know where it was installed. What's it installed on?

Steve: Had a hard time getting out of high school with those multiple choice tests that, you know...

Leo: Oh, lord.

Steve: Installed on the ceiling. Okay, good.

Leo: Okay.

Steve: Windows 11, anyone? Yesterday morning I received the following email from a podcast listener and happens to be a SpinRite 6.1 tester. Jeremy wrote from Texas: "Hi. FYI, about Wednesday of last week, Windows 10 Update offered to 'update' me to Windows 11, even though my HP AIO" - that's an all-in-one - "does NOT" - he has in caps - "qualify because it does not have the latest boot security level. I immediately switched over to GRC.com and downloaded InControl and set my PC to stay on Windows 10 22H2 with security updates. When I reloaded Windows Update, Windows 11 was no longer being offered.

"I thought I was 'safe' from Windows 11," he writes, "since I didn't qualify, but I guess not. Just wanted to let you know to combine my efforts with others, if any. It's a reminder to be proactive rather than rely on assumptions. I've invested a lot of muscle memory in Windows 10 and don't think Windows 11 will offer me much. I don't know if it was a momentary glitch in Win Update or a real offer, since I got InControl immediately." And he signed off "Regular listener and 6.1 tester, Jeremy in Texas."

Well, also yesterday came the news that Microsoft had fixed a bug that was responsible for causing upgrade offers to unsupported PCs. Apparently, the issue came to light last Thursday for Microsoft, and it was quickly resolved, and the fix was then pushed out to affected devices over this past weekend. And this isn't the first time this has happened. Windows 11 22H2 was previously offered to Windows 11 Insiders in the Release Preview channel even when they were using ineligible devices. So it was after Microsoft was aware of this, but before they had pushed the update that Windows users were reporting, many more than just Jeremy in Texas, via Reddit and Twitter, that their unsupported devices were, to their surprise, suddenly being offered Windows 11 upgrades.

So Microsoft confessed this, and they said: "Some hardware-ineligible Windows 10 and Windows 11 version 21H2 devices were offered an inaccurate upgrade to Windows 11. These ineligible devices did not meet the minimum requirements to run Windows 11. Devices that experienced this issue were not able to complete the upgrade installation process." So apparently some users with Windows 10 and even some who were using Windows 11 21H2 were surprised and, I presume, delighted by this news, whereas Jeremy and I would have both been horrified by it. The impacted devices included those running Windows 11 21H2, Windows 10 21H2, and Windows 10 20H2. And, as Microsoft indicated, some later portion of the upgrade process apparently recognized the mistake that had been made and aborted the upgrade.

So, and remember that last month Microsoft announced that it had started a forced rollout of Windows 11 22H2, which is also known as the Windows 11 2022 Update, to systems running Windows 11 21H2, which will be approaching their end-of-support date. Later this year on October 10th, actually, is when it officially ends in of course 2023. And this automated feature update rollout phase came after the Windows 11 2022 update also became available for broad deployment the same day to users with ineligible devices via Windows Update.

Now, as it happens, the little Windows 10 machine I use weekly for zooming this podcast to all of our listeners and Leo at TWiT and everything, it's recently been bugging me, I think the last maybe two or three weeks, to upgrade it to Windows 11. Since I only turn it on before the podcast and then off immediately afterward, I just say no. But I hadn't gotten around...

Leo: You know, there's a wonderful program from a guy named Steve Gibson called InControl. You don't use that, huh?

Steve: Are you reading ahead?

Leo: Oh, no, I'm not, okay. I'd like to be surprised.

Steve: I hadn't gotten around to running GRC's InControl utility.

Leo: Oh, okay.

Steve: What do you know? To tweak the registry to get Microsoft to leave me alone. Now, recall that InControl is the spiritual successor to Never10. And Leo, I still remember your laughter when you first heard that name, Never10. That's right.

Leo: Never11 in this case.

Steve: Eff off Microsoft. Anyway, the expectation was that Never10 would be all that was ever needed.

Leo: Yeah.

Steve: But then Microsoft decided that Windows 10 would not be the last Windows after all. So there was some tendency, you know, rather than create - I thought about creating Never11. That's what people were asking for. But then I'd have to do Never12 and Never13. I wonder if they would do Windows 13? I'll bet not. Anyway, then I'd have to keep changing the name. So instead I switched to InControl, which now no longer at least has to change its name. As long as they leave this facility in place, I won't have to change it at all.

So anyway, the point is that this story from Jeremy and the fact that I kept just saying "no, no, no" every week finally prompted me to take action. And it worked beautifully. What do you know? Initially it was showing the upgrade offer up at the top of the Windows Update screen. So I ran InControl and told it that I wanted to stay put where I was at Windows 10 22H2. So I clicked the button. It tweaked the registry, and that was that. Then I reran Windows Update, and the offer for Windows 11 was still there since it hadn't refreshed.

So I asked Windows Update to re-check for any updates, and that refreshed the screen, and the Windows 11 offer disappeared, and I received what I remembered now when I was originally testing this thing, the little red asterisked notice up at the top of Windows Update which says, asterisk in red, the whole thing's in red: "*Some settings are managed by your organization." Which is, you know, yes, perfect. Leave me alone.

So what we're essentially telling it is that the higher-ups have decided that they're going to be in charge of upgrading, so Microsoft shouldn't be bothering us minions with any of

that. So anyway, I just wanted to remind everybody that that exists because as Microsoft has said, they're going to get more insistent about this; and no, thank you.

Okay. This next item is not security related at all, but something, some news popped up when I was looking for security stuff that I just wanted to make sure everyone was at least aware of. I also have no sense for how large the audience for this might be, but the facility is one of my most-used keystroke combinations. It is "Pasting from the clipboard without formatting." I'm a big user of copy-and-paste for moving things around. But I almost never want to also copy and paste any of the text formatting of the source text which may be present when I move it somewhere else to paste it. You know, it's quite annoying when I paste something, and it jumps into the appearance that it had, typically not even correctly, when all I want is the textual content itself. I want to lose the text metadata.

So finally, a couple of years ago, after several years of annoyance, and I was even doing things like using Windows Notepad as an intermediate stop where in order to force the dropping of the formatting I would open Notepad, copy something that had formatting, paste it into Notepad which can't hold formatting, so it would force it to be lost. And then I would copy that again and paste it into its destination. So I thought, somebody must have fixed this problem.

So I went googling, and I found a slick little utility that I've been using ever since, known as PureText. And that's all it does. It sits in my system tray. It allows you to set any combination of shifts and an action key. I use CTRL-ALT-V rather than just CTRL-V, to perform a non-formatted textual clipboard paste. It's beautiful. Also you can assign a sound if you want, and I know lots of people don't like sounds. I love sounds. So it makes a nice little clunk sound to just confirm that I've got what I wanted. Anyway, I'm mentioning this because Windows Power Tools will, completely separately from that, soon be getting a new module called "Paste as Plain Text." In other words, I'm not the only one, or neither was the guy Steve Miller who wrote this thing, the one I'm using.

As I said, I'm using PureText, both on my Windows 11 and my Windows 10 machines. But Power Tools won't run under Windows 7. You know, you've got to use the Windows Store, and they won't let you do that from Windows 7. So I'll be sticking with PureText. But for what it's worth, if anyone else has this problem, Windows Power Tools will soon be getting Paste as Plain Text. But SteveMiller.net is the site, and this guy's been writing code since Windows 95. He's got stuff dating from then, still some of it useful. So you might want to just check it out anyway: SteveMiller.net.

Okay. Nearly a year ago, last April, The Verge carried the news "Microsoft is adding a free built-in VPN to its Edge browser," with the subhead "Edge Secure Network" - as it's being called, actually now they're calling it Microsoft Edge Secure Network - "will roll out," they said, "as a part of a security upgrade." They didn't say when, and it took like 10 months. But it now appears to finally be happening. It's going to be becoming available.

Two days ago, on Sunday, BleepingComputer posted: "Microsoft Edge's built-in VPN functionality could soon begin rolling out to all users in the stable channel, with some users already getting access to the feature." And they linked to Microsoft's announcement of this. It was the original posting. And one thing's not clear about, from a UI standpoint, what'll be changing. But so here's what we know about what Microsoft is explaining that they'll be bringing to the Microsoft Edge Secure Network.

They said: "Encrypts your connection. Encrypts your Internet connection to help protect your data from online threats like hackers. When using Microsoft Edge Secure network, your data is routed from Edge through an encrypted tunnel to create a secure connection, even when using a non-secure URL that starts with HTTP. This makes it

harder for hackers to access your browsing data on a shared public WiFi network." We know of course that's true, and that's probably the strongest use case for using a VPN, if you're not using one to dial into your corporate network.

Second feature: "Helps prevent online tracking." Right. "By encrypting your web traffic directly from Microsoft Edge, we help prevent your Internet service provider from collecting your browsing data like details about which websites you visit."

Third: "Keeps your location private. Online entities can use your location and IP address for profiling and sending you targeted ads. Microsoft Edge Secure Network lets you browse with a virtual IP address that masks your IP and replaces your geolocation with a similar regional address to make it more difficult for online trackers to follow you as you browse."

And finally: "Is free to use. Get 1GB of free data every month when you sign into Microsoft Edge with your Microsoft Account. A few early adopters will be in a data upgrade trial. At the end of their 30-day trial period, the experience will reflect the normal VPN gigabyte limits." I don't know what that last sentence meant.

Okay. But not surprisingly, my Edge browser doesn't have it yet. You know, I'm not in any advanced bleeding-edge mode. But under Edge's main menu near the menu's bottom, you'll find current entries for "Read aloud" and then "More tools." And assuming that Edge's UI hasn't changed since this posting was last updated, and the rest of it does look the same, there will appear a new "Secure network" entry in between "Read aloud" and "More tools."

The other piece of interesting news is that this is being done in affiliation with our friends at Cloudflare. Microsoft wrote: "Microsoft Edge Secure Network is a service provided in partnership with Cloudflare. Cloudflare is committed to privacy and collects a limited amount of diagnostic and support data acting as Microsoft's data subprocessor in order to provide the services. Cloudflare permanently deletes the diagnostic and support data collected every 25 hours." Now, they didn't say that Microsoft doesn't collect it and won't retain it. We don't know. But they are saying Cloudflare's not keeping it. And presumably Microsoft does not either.

They said: "To provide access, we store minimal support data and access tokens which are only retained for the duration of the required service window. A Microsoft account is required to access Microsoft Edge Secure Network and is retained to keep track of the amount of Microsoft Edge Secure Network data that is used each month. This data retention is necessary to provide 1GB of free Microsoft Edge Secure Network service and to indicate when the data limit has been reached."

So I don't really have any calibration on how quickly 1GB will be consumed, but that doesn't sound like much data for a month. I checked my phone, which I don't use very much for any heavy work because I'm always sitting in front of a computer. And I have the "small" Verizon plan which is limited to 2GB a month. As I said, I'm not doing much with my phone. Turns out I've used than 0.3GB per month for the past three months. So I'm not a heavy data user on my phone. I expect that this might be something that's used sparingly and only when necessary. And, you know, you're not going to watch movies over your 1GB of free Microsoft Edge VPN data.

Their UI does have a "bytes used so far this month" meter, so it'll be possible to track one's usage and get a sense for how it's going and whether you need to scale back and so forth. Anyway, overall this seems like a useful and welcome feature. It's limited, but it's free, and you just need to be logged into Edge or Microsoft through Edge, which I imagine Edge users would be. So it'll be there in a pinch.

Okay. LastPass Incident Update. Yesterday, LastPass provided by far more detail about that second more devastating attack that they suffered. And that's of course the one that inspired the Leaving LastPass podcast and the one that followed that was just titled the numeral "1" when we found out that that's what iteration counts in some cases had been left at. And I have to admit that the forensics which were presented were impressive. This doesn't forgive them in any way from screwing up in the several other ways that we know they did. But as far as forensics examinations go, it's impressive. It's easy to tell a story in retrospect. But as I'm describing what they have determined actually happened, imagine figuring this out. That's, again, as I said, it's impressive.

So I have a link to the incident details. I'm not going to cover all of it because it goes into way more detail than we need. But what they wrote was, LastPass has now learned and explained to us, they said: "To access the cloud-based storage resources, notably S3 buckets which are protected with encryption, the threat actor needed to obtain AWS Access Keys and the LastPass-generated decryption keys. The encrypted cloud-based storage devices house backups of LastPass customer and encrypted vault data." Right, and that's what got away from them.

They wrote: "As mentioned in the first incident summary, certain LastPass credentials stolen during the first attack were encrypted, and the threat actor did not have access to the decryption keys, which could only be retrieved from two locations: First, a segregated and secured implementation of an orchestration platform and key-value store used to coordinate backups of LastPass development and production environments with various cloud-based storage resources." That's the first place where the keys were. "Or a highly restricted set of shared folders in a LastPass password manager vault that are used by DevOps engineers to perform administrative duties in these environments."

Okay. They said: "Due to the security controls protecting and securing the on-premises data center installations of LastPass production, the threat actor targeted one of the four DevOps engineers who had access to the decryption keys needed to access the cloud storage service. This was accomplished by targeting the DevOps engineer's home computer and exploiting a vulnerable third-party media software package, which enabled remote code execution capability and allowed the threat actor to implant keylogger malware. The threat actor was able to capture the employee's master password as it was entered, after the employee authenticated with multifactor authentication, and gain access to the DevOps engineer's LastPass corporate vault.

"The threat actor then exported the native corporate vault entries and content of shared folders, which contained encrypted secure notes with access and decryption keys needed to access the AWS S3 LastPass production backups, other cloud-based storage resources, and some related critical database backups." So yes, Leo, as you gasped, this was quite an attack. I mean...

> **Leo:** And I saw a rumor that it was Plex, that it was a flaw, or they thought it might be a flaw in Plex, which this engineer had been running. So this was a very sophisticated attack. I mean, not just targeted, but they were rooting around in this guy's machine and able to find another flaw.

Steve: Yup. And of course Plex is a media software package, so that tracks with that rumor that you heard.

> **Leo:** Yeah, yeah. And it was online because people often put their Plex servers on the public Internet. So that's, yeah, therein lies your problem; right?

Steve: So as I'm reading this, I'm thinking, wow, no one wants these guys on their tail.

Leo: No kidding, yeah.

Steve: Because wow. Okay. So listen to the steps LastPass has since taken in an effort to recover from this attack.

Leo: Well, they should get Kolide is the first thing they should do.

Steve: Yeah.

Leo: Okay. Go ahead.

Steve: Yeah. Remember how last week I was talking about how difficult it would be to ever be able to trust anything ever again.

Leo: Yeah.

Steve: So they wrote...

Leo: Zero trust.

Steve: Uh-huh. "As we progress through incident response, and as part of our ongoing containment, eradication, and recovery activities related to the second incident, we have performed the following actions, with additional work currently being accomplished in scoping and planning." In other words, they're not even done yet, as of yesterday.

So they said: "With the assistance of Mandiant, we forensically imaged devices to investigate corporate and personal resources and gather evidence detailing potential threat actor activity. We assisted the DevOps engineer with hardening the security of their home network and personal resources. We enabled Microsoft's conditional access PIN-matching multifactor authentication using an upgrade to the Microsoft Authenticator application which became generally available during the incident. We rotated critical and high-privilege credentials that were known to be available to the threat actor." Well, okay, good.

"We continue to rotate the remaining lower priority items that pose no risk to LastPass or our customers." Again, good. "We began revoking and re-issuing certificates obtained by the threat actor. We analyzed LastPass AWS S3 cloud storage resources and applied or started to apply additional S3 hardening measures. We put in place additional logging and alerting across the Cloud Storage environment with tighter IAM (Identity and Access Management) policies enforced. We deactivated prior development IAM users." In other words, that had not been done before, so good. "We enabled a policy that prevents the creation and use of long-lived development IAM users in the new development environment." So that's good. They're changing policies and tightening security when they looked and realized, oh, crap, we've got a bunch of ex-development users who still have credentials here.

Then they said: "We rotated existing production service IAM user keys, applied tighter IP restrictions, and configured policies to adhere to least privilege. We deleted obsolete service IAM users from the development and production environments. We're enabling IAM resource tagging enforcement on accounts for both users and roles with periodic reporting on non-compliant resources. We rotated critical SAML certificates used for internal and external services. We deleted obsolete/unused SAML certificates used for development, services, or third parties." And again, that hadn't been done previously. Good that they did it. Need to make that a policy.

"We revised our 24/7 threat detection and response coverage, with additional managed and automated services enabled to facilitate appropriate escalation." And, finally, "We developed and enabled custom analytics that can detect ongoing abuse of AWS resources."

So, okay. It's evident that things are way better now than they were before, which, you know, that's always sort of been the double-edged sword of this is do you trust better somebody who's learned from their mistake. There's obviously lots of, if I may use the term, "learnings" here. On the other hand, they were clearly doing some things wrong by policy. And that's difficult to forgive, and that's what finally caused us en masse to leave LastPass.

Leo: Well, and also do you want those four DevOps guys who have the keys to be able to bring those home? I guess you do. Now, this all because of COVID, probably; right?

Steve: Yeah. And he was probably VPNing in. One of the things I noticed that I commented on here in passing is they said "We rotated existing production service user keys, applied tighter IP restrictions." [Buzzer sound] If they didn't have any, then Russia could have connected to their S3 buckets; right?

Leo: Right, right.

Steve: So it makes absolute sense to allow, first of all, the IP of somebody VPNing will be the same as the corporate IP network.

Leo: Right.

Steve: So if you weren't VPNing, your IP would be at home. But at least you'd be in the vicinity. So again, the problem we've always had is this tendency not to want support calls; right? It's like, no, we don't want tech support calls. So if we don't default to things being open, then we're going to get complaints when things don't work.

Leo: Right.

Steve: So we're still sort of in that mode. And also, looking at this list of stuff, one of the bugaboos of evolution over time, and LastPass has been around for a long time, is that things tend to become more complicated over time.

Leo: Yeah.

Steve: This is usually driven by inevitably changing requirements. New systems are added to improve or to enable some new job. But the new system doesn't completely take over for the old one. So that older system needs to still stay around to do some of those few things that the newer system doesn't quite do the same. Then the requirements change again, and some customizations are required. Some glue code is created by somebody who then later quits and takes his notes and knowledge with him. Now no one wants to touch that weird box in the corner since no one's quite sure how it works.

And that's the way this, I mean, this actually happens in the real world. And Leo, I've heard you at the beginning of some of the podcasts recently, like trying to figure out how to - I don't know what it is that's going on over there. But, you know...

Leo: Oh, boy. I don't either.

Steve: Exactly. That's what happens. You know?

Leo: I used to know how everything worked here. Now I know nothing.

Steve: Anyone who's been working within a complex environment with many players and constant time pressure, where needs are dynamically changing, will probably be able to relate to this sort of mess that winds up evolving from what was originally a simple solution. And so my point is, in the context of security, this sort of creeping, evolving complexity makes both keeping things truly secure and recovering rapidly from an incident, if one happens, much more difficult.

And it occurred to me that there really needs to be somebody who is assigned the task of stepping back from the day-to-day fray to take sort of a holistic view of an enterprise's systems and be constantly working to reintegrate the inherently disintegrating systems that just naturally form. Keeping things as simple as possible has tremendous benefits for an organization, and in a sufficiently large organization I think it really ought to be a job title. There ought to be like a job title like, I don't know, Holistic System Reintegrator or something. He could, like, not shave and dress funny because that would sort of fit the title. But really, somebody who's working against these otherwise sort of natural forces of entropy which tend to disintegrate things over time.

Okay. Exactly as we predicted, three days ago BBC News headlined their coverage "Signal would 'walk' from UK if Online Safety Bill undermined encryption," with the subhead "The message-encrypting app Signal has said it would stop providing services in the UK if a new law undermined encryption." Signal's president, Meredith Whittaker, told the BBC that if they were forced to weaken the privacy of their messaging system under the Online Safety Bill, the organization "would absolutely 100% walk."

And of course the government said that its proposal is not a "ban on end-to-end encryption." But the bill which was introduced by Boris Johnson is currently going through Parliament. And as we recently covered in detail, under the revisions proposed by this new legislation, companies would be required to scan messages on encrypted apps for child sexual abuse material, language suggestive of "grooming," or terrorism content.

WhatsApp previously told the BBC that it would also refuse to lower its "security," and I put that in air quotes, for any government. In the case of WhatsApp, the question might come down to the definition of the term "security." But the folks behind Signal are likely to be far more clear about that. The BBC's coverage reminds us that the government and prominent child protection charities have long argued that encryption hinders efforts to combat online child abuse, which they say is a growing problem.

The UK's Home Office said in a statement: "It is important that technology companies make every effort to ensure that their platforms do not become a breeding ground for pedophiles." The Home Office added: "The Online Safety Bill does not represent a ban on end-to-end encryption, but makes clear that technological changes should not be implemented in a way that diminishes public safety, especially the safety of children online."

Leo: Children. Think of the children.

Steve: That's right. "It is not a choice between privacy or child safety. We can and we must have both." Right. We can, because we say we can. We're willing to go as far as to change the definitions of words in order to have both the safety of our children and total privacy for everyone, even though we may need to change the meaning of "total," but just a little bit. That pesky math is so annoyingly absolute. After all, we create laws. That's what we do. Unfortunately, not the laws of nature or of mathematics. But still, this is what we want, and we're used to getting our way.

So the UK's child protection charity, the NSPCC, said in reaction to Signal's announcement: "Tech companies should be required to disrupt the abuse that is occurring at record levels on their platforms, including in private messaging and end-to-end encrypted environments." But the digital rights campaigners on the other side, the Open Rights Group, said this highlighted how the bill threatened to "undermine our right to communicate securely and privately."

Signal's Ms. Whittaker said back doors to enable the scanning of private messages would be exploited by malignant state actors and create a way for criminals to access these systems. When asked if the Online Safety Bill would jeopardize Signal's ability to offer a service in the UK, she told the BBC: "It could, and we would absolutely 100% walk away rather than ever undermine the trust that people place in us to provide a truly private means of communication."

Leo: Good. Right on.

Steve: She said: "We have never weakened our privacy promises, and we never will." Period.

Leo: Good.

Steve: Yup. Matthew Hodgson, chief executive of Element, a British secure communications company, said that the threat of mandated scanning alone would cost him clients. He argued that customers would assume any secure communication product that came out of the UK would necessarily have to have backdoors in order to allow for illegal content to be scanned. Matthew added that it could also result in a very surreal situation where a government bill might undermine security guarantees given to

customers at the Ministry of Defense and other sensitive areas of government. He said that his firm might have to cease offering such services. And that raises a great point. Would the most sensitive users within the government also be consenting to having all of their communications intercepted, scanned, and possibly forwarded to a central clearinghouse for human oversight? That seems unlikely.

As for child safety, Signal's Ms. Whittaker said: "There's no-one who doesn't want to protect children. Some of the stories that are invoked have been harrowing." When asked how she would respond to arguments that encryption protects abusers, Ms. Whittaker pointed to a paper by Professor Ross Anderson, which argued for better funding of services working in child protection and warned that "the idea that complex social problems are amenable to cheap technical solutions is the siren song of the software salesman."

Leo: Wow. Wow.

Steve: Yeah. "The idea that complex social problems are amenable to cheap technical solutions is the siren song of the software salesman." So there's no question that the issue of child safety is real. But terrorism content was also mentioned. And doesn't everyone also appreciate that no government, no matter how respectful of its citizens' inherent and often constitutionally guaranteed privacy rights is comfortable with not having some capability for oversight over its citizenry when it believes that such might be needed. As I've noted of our own constitutional government in the U.S., the Constitution's guarantee for privacy is conditional. Courts are able to issue search warrants when presented with probable cause.

We've been watching the approach of this slow-motion collision for years. So it's going to be very interesting to see how this all shakes out. We know that Signal and Telegram and Threema will not capitulate. There's just no way they would. But it's difficult to see how Meta's various services, Google with Android, and Apple may not be forced to go along rather than lose access to those huge markets. Apple already demonstrated a willingness to find some compromise by performing that local fuzzy hash scanning. Of course the backlash from that was epic.

Ultimately, I think governments are probably going to win this legal battle since they're the ones who write the laws, and thus it's possible for them to delineate what behavior is legal and what is not. At that point, any use of fully secure end-to-end encrypted solutions will be outlawed, at which point, as we know, only outlaws will be using them.

Leo: So download it now so you, too, can be an outlaw.

Steve: That's right. Get your soon-to-be-illegal-to-use software.

Leo: Yeah, yeah.

Steve: I do know some people, actually, who have purchased guns because they were worried that they were soon going to be outlawed. It's like, no, no, you don't have to worry about that in the U.S.

Leo: Not in the U.S., baby.

Steve: Okay. Okay, more PyPI troubles. Remember last week how I felt that I needed to at least mention the continuous background of ongoing attacks on open source repositories and registries? Well, last Friday the security firm Phylum posted some news regarding the PyPI, you know, the Python Package Index. Their posting's headline was "Phylum Discovers Aggressive Attack on PyPI Attempting to Deliver a Rust Executable."

They wrote: "On the morning of February 23rd, 2023, Phylum's automated risk detection platform started lighting up with another series of strange publications on PyPI. After digging into it, we were able to link it to another smaller campaign from January, last month. First," they said, "we can confirm that this is an ongoing attack. As we worked on this write-up we saw the list of packages published go from a few dozen to over 500."

Leo: Ugh.

Steve: I know. "The most recent packages appear to be getting published at around one every four to eight seconds, so we suspect that this may continue for some time. You can look at the Package Publication section at the bottom of this post to see the packages we've seen. As of the publication of this post we've already seen 1,138 malicious packages published."

Anyway, they go on to explain in detail the nature of the malware. The short version is that the malicious packages connect to a Dropbox account to download and install a Rust-based malware strain. Phylum says the attacker appears to be the same group that was previously spotted by Fortinet and ReversingLabs the week before in a separate, smaller campaign.

You know, as I was saying last week when we talked about this, our open source repositories are now under more or less constant attack. And Leo, as you commented last week, the industry needs to come up with some solution to the poisoning of our open repositories. The only solution I can see is a future where Internet identity and reputation can be rigorously established and verified. And I know people love the freedom of using synthetic online identities and monikers. And I think that's 100% fine, so long as they don't also want the benefits that accrue from being known and trusted. At least at the moment it's unclear how it's possible to have both. At the moment we have neither. Again, I just, you know, we want this stuff to be open, but why we can't have nice things.

In the interest of giving credit where it's due, the often-in-the-dog-house Taiwanese hardware vendor QNAP, right, the manufacturer of the always apparently in trouble NAS devices which are exposed to the Internet, on Friday announced the launch of its own bug bounty program. Yay. Vulnerabilities relating to QNAP operating systems, applications, and cloud services are all in scope, and rewards can go up to U.S. $20,000.

They do still need to find some way of keeping their devices patched when problems are found; but as we've just seen with VMware, they're not alone in having that problem to solve. I think that tomorrow's systems, that is, globally, one way or another are all going to have to phone home just as our consumer operating systems have all been doing for some time. And there will either have to be an autonomous upgrade and reboot facility, or some reliable notification path to the device's administrators. This problem needs to get solved. But finding and fixing those problems comes first; and QNAP's bounty, while not huge, is a clear step in the right direction. And Leo, let's talk about the Club for a minute.

Leo: I can tell.

Steve: Yeah, I'm going to take a break.

Leo: I know when you get a little thirsty, yeah.

Steve: And then we're going to talk about an unbelievable bug which we have found in motherboard BIOSes, which is going to necessitate a change in SpinRite. And then we're going to talk about what the NSA tells us we should do at home.

Leo: Okay. Okay. Somebody's got to pay for Steve's hydration. Go right ahead.

Steve: So I haven't talked about SpinRite at all for several weeks because I didn't have any significant news to share. But as a result of the work and discoveries over the past week, I have news today which will more than make up for my previous weeks' silence. A major mystery which had been stymieing the project for weeks has been solved. For the past many weeks I've been tracking down, as I have mentioned last time I talked about SpinRite, various sources of SpinRite crashes. I've discovered various sorts of misbehavior from DOS and motherboard BIOSes. They've been altering values in registers that they assume wrongly that others won't be using. You know, that's against every rule of good citizenship. You always put things back the way you found it when you're done using them. But in this case these things weren't.

The only way I can explain this behavior is that they're using, for example, the upper half of 32-bit registers, figuring that in a 16-bit environment no 16-bit app would notice. But SpinRite is now largely 32-bit code, and it makes constant use of all of those extra 32-bit resources. And in other cases they're just not bothering to preserve any part of a register. So a lot of the work I've been doing recently has been defensive computing. Someone suggested that term in our development group; and I thought, well, maybe survival computing is the better term because, I mean, I have to do it in order to keep SpinRite running.

So SpinRite already works without trouble for nearly everyone. But for those for whom it does not work, they have my attention because I'm never sure whose fault the problem is, and I need at least to determine that. Is it something I'm doing that I need to fix? Or is it something outside of my control? And there's no reason SpinRite shouldn't be able to protect itself from anything that a system might throw at it, although that was recently challenged.

One Canadian SpinRite tester, Andre, was able to get SpinRite to crash for him reliably. He had a system with a couple of internal drives and a 160GB USB drive connected. The USB drive was being marked RED by SpinRite, which is a new feature. There's a whole, I mean, SpinRite 6 users are going to - they'll recognize it; but boy, I mean, the last several years have really changed SpinRite, moving it from 6 to 6.1.

Anyway, the USB drive he had was being marked RED, meaning that during its initial appraisal of the drive, SpinRite had found something that wasn't right. That's something that we've been seeing with drives that are quite near death. When the user attempts to select the drive for use, they'll receive a pop-up explanation of exactly what's wrong and, when possible, be given the option to proceed to use that drive anyway.

But that 160GB drive was being marked RED, that wasn't the problem. The problem was that shortly after enumerating those three drives, SpinRite would intercept its own attempt to execute an illegal x86 processor instruction, an illegal opcode. That should never happen. I don't recall what made me suspicious, but I first asked Andre just to try unplugging that damaged USB drive. And sure enough, no crash. Then, with that drive reattached, I provided Andre with an old-school DOS utility called "eatmem." Eatmem simply consumes some amount of RAM memory and then drops back to DOS. Back in the day, this was used to stress-test programs by subjecting them to limited memory situations. But it also has the side effect of changing the location in RAM where DOS will load subsequent programs, since it eats memory from the bottom up.

And sure enough, by "eating" various amounts of memory before running SpinRite, SpinRite's crashing behavior would not occur, or it would occur differently, even with that 160GB RED marked drive connected. Around this same time, one of our SpinRite testing participants, a guy named Paul Farrer, who also knows how to write DOS programs, was experimenting on his own with a similar crash that he and another user were both seeing. That other user had attached a 1TB USB drive to his machine, and it was crashing SpinRite when he tried to run it.

Paul hypothesized that the trouble might be caused by SpinRite attempting to read above the 137GB region of a drive. I have no idea how that occurred to him, but it turned out to be prescient. 137GB is one of those many size limitations that we as an industry were constantly plagued with during the PC industry's early growth. Over and over and over, I mean, it's almost comical in retrospect, we kept outgrowing every upper limit that we assumed would never be exceeded.

The classic story was that the 16-bit Apple II had a maximum of 64KB of main RAM memory - yes, 64K of main RAM memory. And when the IBM PC came out with its initial maximum of 640 kilobytes - so exactly 10 times as much as the Apple II - the story was that during a trade show in 1981, Bill Gates said: "We'll never need more than 640K." Now, today, Bill doesn't recall ever having said that. But it's apocryphal in the industry. Whether or not he did, although today it may seem ludicrous, I can easily imagine, having been active in the PC industry back at the time, that it was seemingly reasonable to say that. And that's the point. These were always reasonable-seeming limitations because none of us who were in the middle of this could have foreseen what has happened since.

The early IDE drives had sizes in the hundreds of megabytes or maybe a few gigabytes if, you know, like you could get some of those, and they were really expensive, like a 4GB drive. So the designers of those drives repurposed the addressing bits which had been used by the original cylinders, heads, and sectors registers to scrounge up a total of 28 bits that they could use to linearly address the sectors on an IDE drive. This was called LBA for Linear Block Addressing. The use of those 28 bits to address sectors meant that a drive could have at most 2^28 total physical sectors. Since sectors were 512 bytes each, that meant that the maximum size of those 28-bit LBA drives would be 137GB. But back then, 137GB, no one would EVER be able to create a drive that large, let alone have that much data that needed to be stored. I mean, come on. That's 137 billion with a "B." Whoops.

So Paul's intuition was that SpinRite's code was somehow being corrupted when it attempted to access sectors at the end of the drive, if a drive was larger than 137GB, like that 160GB drive that Andre had was. And accessing sectors at the end of the drive is one of the things that SpinRite does when it's sizing up a drive before listing it for use by its user. So Paul and I each independently wrote testing utilities to better understand what was going on. And what we discovered over this past weekend was a bit astonishing.

I ended up writing two utilities. The first one was called "BIOSTEST." BIOSTEST first filled all of the system's main RAM memory that wasn't already in use by DOS and the BIOS and buffers and the program itself, with a deterministic pseudorandom pattern. Then it simply used the system's BIOS to read several sectors from the front of the drive and from the end of the drive. After each test read it would rescan all of main memory looking for the first mismatching of data. It was looking to see whether reading any sector from a drive through the BIOS would cause the BIOS to alter main memory. And sure enough, after reading the last sectors of larger than 137GB USB drives on some motherboards, it found main memory mismatches. Reading sectors from the front of the drives never caused any problems. But reading any sectors whose linear address had more than 28 bits would actually damage the data stored in main memory. The USB support in those BIOSes was seriously buggy.

During what SpinRite calls "drive discovery and enumeration," it goes out to the very end of every drive it can locate to perform reading and writing confirmation and confidence testing. It's safest to do that out there, out at the end, because the very ends of drives are usually empty and don't contain any user data. In fact, partitions don't tend to go all the way out to the end. They have like a shorter wrap factor to be aligned. But when SpinRite was doing that on USB-connected drives on motherboards with those broken USB BIOSes, bugs in the BIOS were blasting SpinRite's code, which was then causing it to crash.

Whereas BIOSTEST checked a few sectors at the front and back of every BIOS drive, the second utility I wrote, BIOSSCAN, read every sector of a user-specified BIOS drive. It also filled memory with a pseudorandom test pattern, then it would scan the entire drive from front to back, while rescanning main memory. Anyway, I have a bunch more of this story in the show notes. I've taken up enough of everyone's time. We found a bad bug that some BIOSes have which occurs when an attempt to read the end of drives larger than 137GB happens.

So we've fixed the problem. We're going to clamp SpinRite's access at 137GB for USB-connected drives until we get to SpinRite 7 because it is not safe for SpinRite to go any further than that. There's no reliable way to test whether a BIOS is buggy or not. They're scattered around the SpinRite users that we have, and this is even though it's a large bunch of users, it's a small sample size. So we're going to get SpinRite 6.1 finished. It's one less source of crashing now exists. And the first thing I'm adding to SpinRite 7 is native support for USB host controllers, which will then release this temporary limit on USB drive size and also dramatically speed up SpinRite's ability to use USB drives in the same way that it has sped up SATA and parallel drives, running them now like at their absolute maximum speed they're able to function.

So it was an interesting adventure. With the help of lots of testing we were able to track down a bizarre problem which has gone unseen in the industry. One person did find some reference to a boot manager saying that you should use it at the beginning of drives or not on drives larger than this 137GB because of some bugs that were in BIOSes. But it certainly hasn't been widely known.

I did want to mention an interesting programmer-oriented Humble Book Bundle. We've talked about those in the past. And in fact I found that, because it's got a crazy long URL, I assigned it one of GRC's shortcuts, and I decided to use the shortcut "bundle," B-U-N-D-L-E, which had been taken. I used it before in 2021 on a different Humble Bundle. So we're reusing the same shortcut. It's an 18-book Humble Bundle of programmer-related books.

So anyway, take a look at it. There's a bunch of programming books by Randall Hyde, who's legendary in programming circles, and a bunch more. So grc.sc/bundle, B-U-N-D-L-E, if you are interested. And Leo and I were talking about this before we began

recording. They're not absolutely amazing books, but there are 18 of them. You pay what you want for them. The money goes to support the EFF. And there are some good ones in the collection, as is often the case.

Leo: This "Think Like a Programmer" is a classic, and it's got that assembly language book, too, which you like.

Steve: Yes, "The Art of Assembly Language."

Leo: And this, if you want to learn Clojure, everybody loves this book. Although I think this is online for free. This is all No Starch Press. They do great stuff. I really love No Starch Press's stuff. I have a lot of their books.

Steve: Yup, grc.sc/bundle. Okay. So surprisingly, the NSA has offered some home security "Best Practices" advice. Last Wednesday the National Security Agency, our NSA, published an attractive and end-user accessible nine-page PDF loaded with tips for helping to secure a home network environment. And Leo, it's our shortcut of the week, so if you want to scroll through it on the screen while I'm talking about it, grc.sc/912, grc.sc/912.

Anyway, it is really good. And I want to share and comment on some of what the NSA has suggested. Our longtime listeners will feel right at home, so to speak, with everything that the NSA wrote. They led off with three major key points: Upgrade and update all equipment and software regularly, including routing devices; exercise secure habits by backing up your data and disconnecting devices when connections are not needed; and limit administration to the internal network only. These three points are actually pulled out of the larger piece. But, you know, clearly, keep your stuff up to date. Do secure things. And do not open anything to remote administration.

They also said, and I thought this was interesting: "IoT devices on a home network are often overlooked, but also require updates. Enable automatic update functionality when available. If automatic updates are not available, download and install patches and updates from a trusted vendor on a monthly basis." So it's interesting that the NSA, they, too, see the threat posed by our out-of-date or defective IoT devices. Of course the question is often, who are you going to call you to update some random IoT light switch or wall plug? But moving forward, it would be good to see future devices based on open standards and platforms, and for there to be some sort of certification systems in place. We have a long way to go, but such work is underway.

Okay. And this one was interesting. They wrote: "Your Internet Service Provider may provide a modem/router as part of your service contract. To maximize administrative control over the routing and wireless features of your home network, consider using a personally owned routing device that connects to the ISP-provided modem/router, in addition to modem/router features to create a separate wireless network for guests, for network separation from your more trusted and private devices."

Okay. Now, that was a little bit surprising. They're saying even if your Internet service provider offers a modem/router as part of the service package, get your own that you control and manage, and use it to connect your network to the provider's bandwidth. Again, some sound advice. And on router guidance they say: "Your router is the gateway to your home network. Without proper security and patching, it is more likely to be compromised, which can lead to the compromise of other devices on your network, as well. To minimize vulnerabilities and improve security, the routing devices on your home

network should be updated to the latest patches, preferably through automatic updates. These devices should also be replaced when they reach end-of-life for support. This ensures that all devices can continue to be updated and patched as vulnerabilities are discovered."

Okay. How many times have we seen companies explaining that they won't be offering updates to fix known critical remote code execution problems for older devices because they are EOL, you know, end-of-life, so anyone still using those devices is SOL. And we've often seen, like, inventories of these end-of-life devices still out, exposed on the public Internet, and they're never going to get patched. When selecting a router, this suggests an important criteria that's easily overlooked, and that's the active and supported service life that has historically been provided by various competing vendors.

If this criteria were to become a popular advertised selection, it would put more pressure on vendors to keep older devices supported longer, even though it might mean reduced sales in the future due to the longevity of previous products which were still supported and going strong. I don't know that anybody is actually yet buying a new device to replace a working old device simply because the old device is no longer receiving updates. I don't think most people even know whether devices are receiving updates or not. But I appreciate the NSA saying, you know, if a device you've got is so old that it is no longer receiving updates, and if a problem were found that could never be fixed, and if you care about security, these things aren't that expensive anymore.

The NSA also talked about WPA3. We briefly touched on this next-generation Wi-Fi 6 and WPA3 encryption, but we haven't yet given it a deep dive, and it's probably time for us to do so. It's had a somewhat slow liftoff, since the Wi-Fi Alliance's WPA certification process started back in 2018, so between four and five years ago. But Wi-Fi 6 and WPA3-capable devices are here now. So we'll get around soon to doing a podcast.

Here's what the NSA wrote. They said: "To keep your wireless communications confidential, ensure your personal or ISP-provided WAP is capable of WiFi Protected Access 3. If you have devices on your network that do not support WPA3, you can select WPA2/3 instead. This allows newer devices to use the more secure method while still allowing older devices to connect to the network over WPA2.

"When configuring WPA3 or WPA2/3, use a strong passphrase with a minimum length of 20 characters. When available, protected management frames should also be enabled for added security. Most computers and mobile devices now support WPA3 or 2. If you're planning to purchase a new device, ensure that it is WPA3-Personal certified. Change the default service set identifier (SSID) to something unique. Do not hide the SSID as this adds no additional security to the wireless network and may cause compatibility issues." All of that is true. I was very impressed as I was reading through this that, you know, the degree to which the NSA got it.

So as I said, we'll do a Wi-Fi 6 podcast soon. Seeing this next one raised an eyebrow since everyone knows that I worry about the day a widely used IoT device goes rogue. The NSA wrote: "Implement wireless network segmentation." They said: "Leverage network segmentation on your home network to keep your wireless communication secure. At a minimum, your wireless network should be segregated between your primary WiFi, guest WiFi, and IoT network."

Leo: Wow. You've been saying this for years.

Steve: Yup.

Leo: But it's interesting to hear them recommend this to normal people because I don't think normal people know how to do this.

Steve: I know. I know. And they finish, saying: "This segmentation keeps less secure devices from directly communicating with your more secure devices." As we know, I've been promoting multi-NIC routers which are able to do that. And some of the more recent WiFi routers are beginning to offer stronger segmentation options, as well. So I would say that's something to look for when you're shopping for a new WiFi router is check out and see if it offers built-in segmentation. That will make it a lot more accessible to average users.

Leo: What do they call it? Is there a name that they give it?

Steve: That's a good question. I think that the ASUS router I'm using, it definitely has a guest WiFi. But I think it has multiple guest WiFis. And so you could, and you are able to, keep them from talking to each other. So that would mean you could give your second guest WiFi, use that for your IoT stuff.

Leo: Okay. The problem I think a lot of people will have is that they can't then use devices on the main network to control the IoT devices. And you and I and our audience probably knows how to do sophisticated firewall rules to allow that. But you have to know that's really, you know, Black Diamond stuff; right? That's pretty advanced.

Steve: Yeah. Yeah. But again, it was amazing to see the NSA saying this. And our users get a little additional impetus behind that, too.

Leo: Well, and it looks like everything in this NSA document you would agree with 100%; right?

Steve: Yes, absolutely. In fact, there's something I've never recommended that I agree with that we'll get to in a second.

Leo: Oh, oh.

Steve: So what about the presence of personal assistant technologies and worries over eavesdropping? Well, not surprisingly, the NSA is not a big fan of things with microphones.

Leo: Unless it's theirs.

Steve: Unless they're on the other end, exactly. So they wrote: "Be aware that home assistants and smart devices have microphones and are listening to conversations, even when you are not actively engaging with the device. If compromised, the adversary can eavesdrop on conversations. Limit sensitive conversations when you're near baby

monitors, audio recording toys, home assistants, and smart devices. Consider muting their microphones when not in use."

Leo: Wow.

Steve: "For devices with cameras - laptops, monitoring devices, and toys - cover cameras when you're not using them. Disconnect Internet access if a device is not commonly used, but be sure to update it when you do use it." So I got a kick out of that one. And all that security advice falls nicely under the umbrella of generally sound, if maybe a little paranoid, security advice. Following that, under the topic of general security hygiene, they add: "To minimize ransomware risks, back up data on external drives or portable media. Disconnect and securely store external storage when not in use." You know, take it offline.

"Minimize charging mobile devices with computers; use the power adapter instead. Avoid connecting devices to public charging stations. Leave computers in sleep mode to enable downloading and installing updates automatically. Regularly reboot computers to apply the updates. Turn off devices or disconnect their Internet connections when they will not be used for an extended time, such as when going on vacation."

In other words, think security at all times, and try to never take it for granted. It's sort of the broader equivalent of what has happened to email, where it's no longer ever safe, unfortunately, to assume that all email is legitimate and that links can be clicked on without careful scrutiny. It's a sad state, but it's the state we're in.

And everyone knows that I love this one: "Limit administration to the internal network only." They said: "Disable the ability to perform remote administration on the routing device. Only make network configuration changes from within your internal network. Disable" - this is them. "Disable Universal Plug-n-Play. These measures help close holes that may enable an actor to compromise your network." And Leo, I do kind of wonder maybe if they listen to the podcast.

And there was one piece of advice that makes sense, but I have never recommended. They said: "Schedule frequent device reboots." They wrote: "To minimize the threat of non-persistent malicious code on your personally owned device, reboot the device periodically. Malicious implants have been reported to infect home routers without persistence. At a minimum, you should schedule weekly reboots of your routing device" - that's a little often, but okay - "smartphones, and computers. Regular reboots help to remove implants and ensure security." I mean, it's true. They do.

What's interesting about this advice is that, as we know, many forms of malware are RAM resident only. They never write anything to non-volatile media. Some routers are almost never rebooted, so malware authors probably figure that there's no reason to bother writing it to non-volatile memory and arranging to get it to start where it's a little more visible in the startup script. And we know that in well-protected environments, writing to disk can trip all sorts of monitoring alarms. And some malware might want to disappear after a reboot so that its larger network of devices can remain hidden. So if it's not necessary for something to survive a reboot, malware might well choose not to.

Consequently, indeed, a reboot will permanently flush RAM-based malware from the system. Okay. Now, if the way such malware originally got into the system in the first place and then obtained its foothold in RAM, if that's not closed off and resolved, then it might come back before long. But, yeah, reboots are inherently cleansing. I think that was a great point, and it's one I've never talked about before.

So okay. Those were just some of the highlights that I thought were the more interesting, a little bit surprising in some cases, and insightful. But there is much more than those in the nine-page document than what I've just shared, and the entire document is so good that I think everyone listening would benefit from at least scanning and probably also by recommending it to others. It has the additional pedigree of bearing the official seal of the National Security Agency, which might help everyone's non-Security Now! listening friends sit up a little bit and take it seriously. And, as we've seen, it's far from being the typical useless piece of say-nothing bureaucratic nonsense.

As friendly and useful as the document is, its line-and-a-half wrapping around URL is not nearly as friendly. So this week's GRC shortcut is that. You can find it at grc.sc/912, since this is Episode 912. And big props to the NSA for assembling something so useful and largely so actionable. If nothing else, the nature of the recommendations would help someone who doesn't live in the security realm to realize the way security-conscious professionals think. And that would probably be surprising to many people. It's like, wow, you're really that paranoid? I'm sorry.

Leo: Yeah, yeah. I'm going to share this on Ask the Tech Guys. This is really good.

Steve: Yeah. It is a great document.

Leo: Yeah, yeah. And, yeah, I can't think of anything I would disagree with. There are some things that maybe are not easy to implement.

Steve: Yeah, I mean, like going around, like cover up the microphone holes on your various devices.

Leo: Yeah. I mean, baby monitors we know get hacked. But no device from Amazon or Google to my knowledge has ever been hacked. Most people are more afraid of Amazon and Google listening in, which they don't, or Siri. I think that the big three, right, you don't know of any hacks of them, exploits with them?

Steve: No.

Leo: Of course if you had an exploit, you wouldn't tell anyone. That'd be a good nation-state exploit. I know a surprising large number of people that we work with, sophisticated users, will not have these devices in their homes. I have an Echo right there, a Google Assistant right there. I have Siri right here. You've got it in your pocket with your phone. But they just say no, no, no, not going to do it.

Steve: Yeah.

Leo: Very interesting. Great show, as always, Mr. G. Steve Gibson is at GRC.com, the Gibson Research Corporation.