Password Manager Injection Attacks

Description: What happened during Microsoft's recent Windows Endpoint Security Ecosystem Summit? And what, if anything, will probably result? How reliable is ANY form of digital storage when used for long-term archiving? What happened when an illegal Starlink Internet network was set up on a U.S. Navy ship? What's the best solution for securing the Internet-facing "edge" of enterprise networks? GRC has started notifying SpinRite 6 owners about 6.1. What's been learned about the challenge of sending email in 2024? Why might running SpinRite on an SSD cause the SSD to then appear to be running more slowly? Why is true secrecy so difficult to achieve, and how were most password managers leaking some of their secrets?

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-992.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-992-lq.mp3

SHOW TEASE: Coming up on Security Now!, I, Mikah Sargent, am subbing in for Leo Laporte again this week. Steve Gibson kicks things off by talking about what the heck happened at the recent Microsoft Summit, where they, you know, aimed to talk about what went wrong with CrowdStrike. Plus we talk about storage and how, regardless of what you're using to store data, all of it falls to entropy in the end. Then a conversation about a really fascinating investigation regarding Starlink and the U.S. Navy, and a very important conversation about password managers and how many of them are vulnerable to attacks. All of that, plus so much more, coming up on Security Now!.

MIKAH SARGENT: This is Security Now!, Episode 992, recorded Tuesday, September 17th, 2024: Password Manager Injection Attacks.

It's time for Security Now!, the show where we talk to THE Steve Gibson about cybersecurity each week. I am Mikah Sargent, subbing in once more for Leo Laporte. Hello, Steve Gibson, reporting from another secret cave buried deep underground.

Steve Gibson: An undisclosed location, yes.

MIKAH: Good to see you.

Steve: Likewise, Mikah. Great to be with you for our second of two, you standing in for Leo while he's out somewhere.

MIKAH: Gallivanting.

Steve: Gallivanting around the territory.

MIKAH: Yeah.

Steve: Yeah. I'm looking forward to this. I mentioned this last week, that I'd had some outreach from a security researcher who has previously been the source of many of our interesting and many times wacky sort of side-channel leakage issues. So today's podcast, 992 - yes, closing in

on 999, but no longer the end of the road for the podcast - is titled Password Manager Injection Attacks. And of course the idea that password managers could be in any way subject to security trouble gets everybody's attention.

MIKAH: Yeah.

Steve: They actually - they did two papers, one on password manager injection attacks and another on end-to-end encrypted messaging apps, specifically WhatsApp and Signal injection attacks. I'll talk about that a little bit, just to kind of give some context. But because everyone is using password managers, and the last thing you want is them to have a security problem, you know, that gets center stage for us. But we're also going to talk about what happened during Microsoft's recent Windows Endpoint Security Ecosystem Summit, which, you know, was the thing that they did in reaction to the CrowdStrike global outage from a month and a half ago; and what, if anything, will result from that.

Also, and completely sort of off on the side, but of interest to our listeners, is how reliable is any form of digital storage when used for long-term archiving? That's, you know, CDs, tape, hard drives, solid-state memory. It turns out that we've got problems in terms of long-term archiving. And, my god, talk about an explosion in the amount of storage and thus the amount of problem that we have. Also, what happened when an illegal Starlink Internet network was set up on a U.S. Navy ship without permission, which is a big no-no, as we'll see. What's the best solution for securing the Internet-facing edge of enterprise networks? A security magazine reviewed all of the contenders, and I'm pleased that the one that had impressed me previously last Halloween turned out to win. We'll come back to that briefly.

Also I finally started rolling out the notification of SpinRite 6.1 to all of 6.0 owners, and I've learned a lot, which I'll talk a little bit about. Also why - this is actually a common question that I've had as a result of that - might running SpinRite on an SSD appear to make the SSD run more slowly, rather than more quickly. And finally, the moral of the story underlying these password manager injection attacks is why is true secrecy so difficult to achieve, and how were most password managers leaking some of their secrets? So I think a great and interesting podcast for our listeners. And we've got another great Picture of the Week, as well. So I think a lot of fun for everybody.

MIKAH: Absolutely. I am very much, I don't know if looking forward to it is the right word, the stuff involving password managers. But can't wait to hear about it, for sure. All right. Back from the break, and let's kick things off with a wonderful photo.

Steve: Okay. So it didn't occur to me just until I was looking at this picture, as you were talking about our sponsor, that maybe I should have blurred this person's license plate. I didn't take the picture myself, so it's out on the Internet, so it's already out there. But still, you know. Anyway, this is - and he probably doesn't care. He's probably liking the...

MIKAH: Attention, yeah.

Steve: The attention, yes, because clearly he's into attention. Okay. So this starts with a GMC truck which - and this is apparently kind of a techie owner because we notice from his back window in the lower left-hand corner he's got this - it's stenciled Transmission Lineman, as in one of those high-tension tower, like makes your hair stand on edge, you want to be very careful like with every single movement that you make sort of high-tension towers. Anyway, apparently he noted that the dual exhaust ports at the bottom back of his truck are angled such that they actually really do resemble HDMI ports. And so he went to the trouble of labeling them "HDMI 1" and "HDMI 2" on the back of his truck.

And so the truck was parked somewhere, and somebody thought that was kind of cool and took a picture of it. I also noted that, and the caption I gave this picture was "Not only does this truck have dual HDMI outputs, but great signal strength and is fully charged." Because in the upper right-hand corner of the back window we see that he's added four bars and full WiFi strength and 100% charge on his car. So anyway, clearly a techie.

MIKAH: I'm thinking about adding those. Not the HDMI because it wouldn't work for my little Subaru, but I love the WiFi and [crosstalk].

Steve: That's just great, yeah. That's just a kick.

MIKAH: And yes, they do look like HDMI ports.

Steve: Oh, my god.

MIKAH: You know, I'm just imaging the giant HDMI port that some clown would like, get out of the car and run behind and plug it in the back. Oh, just funny.

Steve: Well, and if it were an eCar, then, you know, that would be great if it was the charging port.

MIKAH: Oh, so good, so good.

Steve: Okay. So we recently noted that Microsoft was responding, as we and they knew they must, to what's now being referred to as the Cloudflare Outage. Wait. I wrote Cloudflare. I don't mean that. CrowdStrike. Wow. Sorry, Cloudflare. CrowdStrike Outage. Duh. It was late last night. Although, you know, as we know, the proximate cause of this global meltdown was a bad update to CrowdStrike's kernel-level code, which forced a Windows kernel panic and thus shutdown, you know, the fact that third-party vendors are, first of all, allowed to install their code into the Windows kernel; and, secondly, that Windows lacked any graceful resiliency, as we all painfully saw, that would have allowed it to somehow arrange to get back on its feet using some sort of rollback to the pre-CrowdStrike update. So, you know, this all meant that Microsoft also received some measure of blowback themselves.

So this summit was held last Tuesday. They called it their Windows Endpoint Security Ecosystem Summit. And then just this past Sunday, Microsoft posted about what happened, who attended, and some of what was said. So this is what they shared with us.

They said: "On Tuesday, September 10th, we hosted Windows Endpoint Security Ecosystem Summit. This forum brought together a diverse group of endpoint security vendors and government officials from the U.S. and Europe to discuss strategies for improving resiliency and protecting our mutual customers' critical infrastructure. Although this was not a decision-making meeting" - which, okay, be nice to have some decisions, but no, no. They said: "We believe in the importance of transparency and community engagement. Therefore, we're sharing the key themes and consensus points discussed during the Summit, offering insights into our initial conversations." There is a hand gesture I can make at this point, but this is a podcast for families, so I won't do that.

"We want to thank every one of our Summit attendees for dedicating their time to participating in these meaningful discussions. The CrowdStrike incident in July underscored the responsibility security vendors have to drive both resiliency and agile, adaptive protection. And it was inspiring to see the engagement throughout the event's agenda and activities. Together with our Microsoft Virus Initiative" - they'll refer to that later as MVI - "partners - companies who develop endpoint protection and additional security products for Windows, covering client, server and IoT - we discussed the complexities of the modern security landscape, acknowledging there are no simple solutions." Right? Otherwise we would have done that.

"A key consensus point at the Summit was our endpoint security vendors and our mutual customers' benefit when there are options for" - there is some interesting reading between the lines here; right? - "when there are options for Windows and choices in security products. It was apparent that, given the vast number of endpoint products on the market, we all share a responsibility to enhance resiliency by openly sharing information about how our products function, handle updates, and manage disruptions.

"In the short term, we discussed several opportunities to improve how we support the safety and

resiliency of our mutual customers. First, we spent time going into depth on how we employ Safe Deployment Practices at Microsoft" - in other words, you know, this is what we did, folks. Why aren't you doing that? Anyway, they said, "...and where we could create shared best practices as a community, including sharing data, tools, and documented processes. We face a common set of challenges in safely rolling out updates to the large Windows ecosystem, from deciding how to do measured rollouts with a diverse set of endpoints to being able to pause or rollback if needed.

"A core SDP" - that's their Safe Deployment Practice. "A core SDP principle is gradual and staged deployment of updates sent to customers. Microsoft Defender for Endpoint publishes SDPs; and many of our ecosystem partners such as Broadcom, Sophos, and Trend Micro have shared how they approach SDPs, as well. This rich discussion at the Summit will continue as a collaborative effort with our MVI partners" - that's the virus initiative partners - "to create a shared set of best practices that we will use as an ecosystem going forward."

So, right. In other words, CrowdStrike, you didn't do that. And I believe the proper term would be "bitch slapped," at this point, for not enforcing any sort of staggered rollout of their updates; right? I mean, that's like the most obvious thing you could have done that would have caught this immediately. Instead, the entire world was hit with this thing at once, and the entire world went down, you know, everyone using CrowdStrike. So this, you know, there's no way to read this from CrowdStrike other than it's suggesting a level of "we can do no wrong" arrogance that did indeed come back to bite them. And, you know, the lack of, like, it's just impossible to justify that in retrospect. There's no good answer to the question, "Why the heck weren't you deploying any sort of staggered rollout?" But, you know, they weren't.

Anyway, Microsoft continues, saying: "Beyond the critical Safe Deployment Practices work, there are several ways we can enhance our support for customers in the near term. Building on the Microsoft Virus Initiative program we have, we discussed how Microsoft and partners can increase testing of critical components, improve joint compatibility testing across diverse configurations, drive better information-sharing on in-development and in-market product health, and increase incident response effectiveness with tighter coordination and recovery procedures. These are a sampling of the topics we plan to make rapid progress on, to improve our collective customers' security and resiliency.

"In addition, our Summit dialogue looked at longer term steps serving resilience and security goals. Here, our conversation explored new platform capabilities Microsoft plans to make available in Windows, building on the security investments we have made in Windows 11. Windows 11's improved security posture and security defaults enable the platform to provide more security capabilities to solution providers outside of kernel mode." And of course that's a key aspect of this; right? The problem is in order for endpoint security to do what it needs to do, because Microsoft has been relatively stingy about what they allow user-mode code to do, it's not able to get the deep access that it needs. Microsoft is saying, well, Windows 11 is better.

Now, my eyebrows went up when I saw that because we're approaching end of life for Windows 10. And this is going to be controversial because nobody wants 11. And I mean, the enterprise has not moved. And so there's a massive install base of 10 that won't be able to have these things because Microsoft is saying, well, it's in 11, so you'd better move. The problem, as we know, is that they've also imposed arbitrary hardware constraints on the upgrade to 11 so that there's a huge install base of hardware that can't upgrade to 11 as it stands now. So this is all - we're going to have some interesting times coming up this year. And I'm glad we're going past 999.

So they said: "Both our customers and ecosystem partners have called on Microsoft to provide additional security capabilities outside of kernel mode which, along with SDP, can be used to create highly available security solutions." In other words, everybody's saying, look, we don't want to be in the kernel either. It terrifies us because of something like this happening. But you're not letting us do what we need to do from user mode.

And I would argue that that's probably never going to happen. Not because Microsoft wouldn't want to, but it's difficult to have user-mode hooks in the kernel that won't dramatically slow down Windows because in order to do what needs to be done, you need to not have user mode kernel mode ring transitions constantly. And so to deeply get into the kernel, that's the only place you can watch everything going on without significantly slowing down Windows.

So, I mean, it really - it's truly a problem, which I understand that Microsoft gets, and it's why they're saying, well, the truth is they can't allow all this to be done from user mode or there'll be constant user-mode and kernel-mode transitions that will dramatically slow things down. It's the reason that GDI, the Graphics Device Interface, was moved into the kernel. We saw huge security implications and compromises as a consequence of that decision. And on one hand Microsoft can be, you know, held responsible for that. The flipside is, back when systems were a lot slower than they are today, they had no choice because the graphics device interface level had to make such deep and frequent access to the kernel. That's why they moved it into the kernel.

Anyway, so they said: "Some of the areas discussed during the Summit include performance needs and challenges" - oh, there it is - "outside of kernel mode; anti-tampering protection for security products" - right, because if you're in user mode you don't have the protections afforded by kernel mode - "security sensor requirements; development and collaboration principles between Microsoft and the ecosystem; and secure-by-design goals for future platforms."

They said: "As a next step, Microsoft will continue to design and develop this new platform capability with input and collaboration from ecosystem partners to achieve the goal of enhanced reliability without sacrificing security." And I'll just note that this is all happening as a consequence of that outage. None of this would be happening if it weren't for that. In other words...

MIKAH: And isn't that - how often does it take something going wrong for people to start doing things correctly?

Steve: Right.

MIKAH: Is that not the crux of so many issues? It's really frustrating that - and it shouldn't have to be that way. It shouldn't have to require things going this poorly to go, oh, right, now's the time. But it's just so built into the way things work.

Steve: Well, and in this case, there had been sort of a - sort of a dtente had been achieved where Microsoft didn't like the fact that they had to open the kernel. They did because the EU forced them to because Windows Defender has access to the kernel. And they're saying, well, if your own endpoint protection technology, if you've given that technology kernel access, then you must open it up to the competition, in the interest of creating competition. And arguably these third-party products like CrowdStrike do a much better job than Microsoft's endpoint protection system does. So customers are getting a better result, at the consequence of this kind of event being possible. And I will actually be surprised if much changes.

I think this is face-saving for Microsoft. They had to do something to respond to the outage. And so, oh, we're going to have a summit, and we're going to get everybody together, and we're going to figure out how to prevent this from happening. The fact is CrowdStrike should have never let this happen because these things like incremental deployment are trivial to do. They're doing it now. I mean, they have said, you know, it's already in place. They will never let this happen again. Well, they shouldn't have let it happen the first time. So there was some arrogance on their part. They've learned their lesson. But the fact is to do what these products need to do, the truth is they have to be in the kernel. They have to be allowed that kind of access because it just - it can't be done up in userland.

Anyway, so they had this first meeting. They quoted a bunch of their partners coming out of this. Adam Bromwich, who's Broadcom's CTO (Chief Technology Officer) and Head of R&D for their Enterprise Security Group, was quoted saying: "Organizations today benefit from a diverse, layered security defense. As a result, industry collaboration is vital to helping organizations stay ahead of persistent threats and remain resilient when unexpected business disruptions occur. As a long-time Microsoft Virus Initiative" - so MVI - "Partner, Broadcom recognizes that working closely with Microsoft and other security vendors not only helps improve our customers' security posture, including endpoint protection, but also the greater global digital ecosystem."

Drew Bagley, the VP, Counsel for Privacy and Cyber Policy from CrowdStrike, the bad guys in this particular stumble that caused this summit to be created, said: "We appreciated the opportunity

to join these important discussions with Microsoft and industry peers on how to best collaborate in building a more resilient and open Windows endpoint security ecosystem that strengthens security for our mutual customers." In other words, he said nothing.

MIKAH: Thank you, because that's what it says to me.

Steve: Yeah, exactly. It was like, okay, fine. We're sorry that we brought the Earth to a standstill. We realize the error of our ways, and we won't let it happen again. And they quoted ESET and SentinelOne and Sophos and Trellix and Trend Micro. You know, Trend Micro: "I applaud Microsoft for opening its doors to continue collaborating with leading endpoint security leaders to make our mutual customers even more cyber resilient. Looking forward to more collaboration." And as I said, my feeling is that this was mostly for show. In fact...

MIKAH: Everybody probably - I'm saying this tongue in cheek. Everybody probably walked away with a nice tote bag filled with lots of Microsoft goodies. It's like, and here's your gift card for this quote, and here's yours for this. It just feels very hollow.

Steve: Yes. Well, and the photo that accompanied this blog posting was a perfect representation of this meeting.

MIKAH: Love the signs.

Steve: Oh, my god. A bunch of executives of various stripes, sitting around a conference table in a stunningly opulent office building conference room setting. You know, there's the requisite whiteboard and a UI projected onto another screen. Then we have the four world time digital clocks which are visible showing the time in London, Moscow, Beijing, and Sydney.

MIKAH: For whatever reason.

Steve: Exactly. For, like, what? And the one laptop that we can see open is distinctly a Mac.

MIKAH: Yup.

Steve: So okay. I don't know what that means. Maybe the guy's installed Windows 11 on his Mac because it runs better than it does on the Surface. Anyway, I think my point is that nothing ever really gets accomplished at these kinds of meetings. This was all just for show, for the government and for Microsoft shareholders. You know, it's okay, you know, we realize we did bad. We're responding to this problem. What's actually going to happen now will be a long, multiyear series of slow, plodding, back-and-forth negotiations where Microsoft will present and may implement some next-generation set of userland hooks for use by their various third-party vendors. The vendors will examine them and explain how what Microsoft is offering still doesn't give them the total freedom that they really want, and I think they can argue they need. Which is only still available through true kernel-level operation.

And so it will go back and forth. Perhaps something will eventually come of it. But, you know, that's far from certain, I think, at this point. So long as this meeting has been held and the parties are now "working on it together," face has been saved, lawsuits will trundle forward, and the vendors will all work harder not to make another similar horrible mistake. As I said, if CrowdStrike had not made this mistake, this would have never happened. It would never happen on its own because things were kind of okay. You know, I mean, the third-party vendors had the deep access they needed. They hadn't brought the world to a standstill. And Microsoft had given only as much as they had to, you know, access and user mode.

My only hope was that since Mark Russinovich, who is truly a serious security kernel-level guy, we all know him from his founding of Sysinternals back in the day, which Microsoft then bought, Mark Russinovich tweeted about this, saying that this really did represent some future hope. So I have some more hope that something might actually change, but nobody should be holding their breath. We know that in retrospect CrowdStrike now realizes it needs to be able to catch anything like this before it is ever rolled out to the entire world. And it's trivial to do. That's what's so mindboggling is it's not like this is rocket science to do an incremental release. Everybody else

does that.

So anyway, we know that they're going to be doing it. They've got that in place already. It also means that everyone else must do the same and never fail at this trivial-to-implement requirement. So, you know, no more cowboy developer jock behavior. The stakes are now far too high.

MIKAH: Amen. Amen. All-righty. Back from the break, and let's talk about how - how it's a good idea, not a good idea, maybe it's safe, maybe it's not. There's a problem with saving stuff.

Steve: Yeah. A listener forwarded this piece from Ars Technica to me, which was titled "Music industry's 1990s hard drives, like all hard disk drives, are dying." And the subhead is "The music industry traded tape for hard drives and got a hard-earned lesson."

MIKAH: Yikes.

Steve: So I'll just share a bit of what they said. They said: "One of the things enterprise storage and destruction company Iron Mountain does is handle the archiving of the media industry's vaults. What it has been seeing lately should be a wake-up call: Roughly one-fifth" - so one out of five - "of the hard disk drives dating back to the '90s it was sent are entirely unreadable."

MIKAH: Wow.

Steve: Yeah, you just think, okay, here's a drive. You guys store it for us. And we may need the data in the future. We'll let you know. And then we'd like it back.

"Music industry publication Mix spoke with the people in charge of backing up the entertainment industry. The resulting tale is part explainer on how music is so complicated to archive now, part warning about everyone's data stored on spinning disks. Robert Koszela, global director for studio growth and strategic initiatives at Iron Mountain, told Mix: 'In our line of work, if we discover an inherent problem with a format, it makes sense to let everyone know. It may sound like a sales pitch, but it's not; it's a call for action.'

"Hard drives gained popularity over spooled magnetic tape as digital audio workstations, mixing and editing software, and the perceived downsides of tape, including deterioration from substrate separation and fire. But hard drives present their own archival problems. Standard hard drives were also not designed for long-term archival use. You can almost never decouple the magnetic disks from the reading hardware inside, so that, if either fails, the whole drive dies.

"There are also general computer storage issues, including the separation of samples and finished tracks, or proprietary file formats requiring archival versions of software. Still, Iron Mountain tells Mix that: 'If the disk platters spin and aren't damaged,' it can access the content. But 'if it spins' is becoming a big question mark. Musicians and studios now digging into their archives to remaster tracks often find that drives, even when stored at industry-standard temperature and humidity, have failed in some way, with no partial recovery option available." So it's completely dead.

"Koszela says: 'It's so sad to see a project come into the studio, a hard drive in a brand-new case with the wrapper and the tags from wherever they bought it still in there. Next to it is a case with the safety drive in it. Everything's in order. And both are bricks.'"

MIKAH: No.

Steve: "Mix's passing along of Iron Mountain's warning hit Hacker News earlier this week, which spurred other tales of faith in the wrong formats. The gist of it: You cannot trust any medium, so you copy important things over and over, into fresh storage. Optical media rots, magnetic media rots and loses magnetic charge, bearings seize, flash storage loses charge, et cetera. Entropy wins, and sometimes much faster than you'd expect.

"There's a discussion of how SSDs are not archival at all; how floppy disk quality varied greatly between the '80s, '90s, and 2000s; how Linear Tape-Open, a format specifically designed for long-term tape storage, loses compatibility over successive generations; how the binder sleeves we put

our CD-Rs and DVD-Rs in have allowed them to bend too much and stop being readable."

MIKAH: Oh.

Steve: I know. One thing after another, one format after another failing. They said: "Knowing that hard drives will eventually fail is nothing new. Ars wrote about the five stages of hard drive death, including denial, back in 2005. Last year, backup company Backblaze shared failure data on specific drives, showing that drives that fail tend to fail within three years, that no drive was totally exempt, and that time does, generally, wear down all drives. Google's server drive data showed in 2007 that hard disk drive failure was mostly unpredictable, and that temperatures were not really the deciding factor. So Iron Mountain's admonition to music companies is yet another warning about something we've already heard. But it's always good to get some new data about just how fragile a good archive really is."

So I can speak for myself. I run a bunch of my own servers, you know, my own actual hardware. I think I have four separate physical servers. Each of the servers has four hard drives. Actually, one of them may still have SSDs, but I'm beginning to swap them out because I have not found them to be more reliable then lower size, smaller size, and by that I mean like 2TB is sort of what I've settled on, spinning drives. Every single one of those is running RAID 6, which is two drives of redundancy. So any two of the four drives could fail, and I lose nothing. And all of the RAIDs, all four of the RAID arrays are being monitored continuously.

And every so often I receive email telling me that one of the drives has failed. So that's okay. Another one could still fail in that group of four, and I'd still lose nothing. So it's not an emergency, but within a day or two I go to Level 3, you know, go through all the security measures, get access to my hardware, pull the dead one out, put a new one in. It spins up, and the RAID rebuilds itself, reestablishing its RAID 6 two-drive failure. And it happens maybe every six months or so. And in fact, right here, this is - I have it because I've been meaning to run SpinRite on it. This is the most recent drive to die. It's a 2TB Seagate Barracuda hard drive. Something about it my system doesn't like. And so I'll run SpinRite.

The time before that this happened, there was - sectors are not actually 512 bytes any longer. They still look like they're 512 bytes because they always were in the past. They're actually 4K. It's much more efficient to have larger physical sectors because you get much more efficiency from error correction, and you just don't need all of the - all of the gaps between sectors take up space. So what I found was there was a single block of what looked like eight logical sectors, but eight 512 byte, which is half a K, so eight half a K sectors is actually one 4K physical sector. It was bad, and the RAID said, okay, I quit. Bad drive.

Anyway, I ran SpinRite on it. It said, oh, you've got a contiguous run of eight sectors which have a problem. SpinRite did what it could to recover it. It rewrote the sector, and it was fine. And that was all that was wrong with the drive. In other words, nothing actually wrong. Just a tiny region, some set of bits that were uncorrectable, and so the RAID said, sorry, this drive's no good. In fact, it was fine. It just needed a little bit of fixing, and then it was good to go.

So there is good reason to believe that performing a periodic rewrite of either magnetic spinning or electrostatic solid-state mass media is an extremely useful thing to do. And you don't need SpinRite to do that, although SpinRite does make that easy, and it provides a great deal of feedback about the state of the drive. And if there is any sort of trouble, it'll fix it for you. But we learned, for example, a few years ago, and it was a surprise at the time, but it's not anymore, that offline SSDs just sitting on a shelf tend to lose their data more rapidly when they're stored at high temperature. Since SSD storage is just about charge leakage, it makes sense that higher temperature would tend to weaken the strength of the dielectric insulation which isolates the charge bits.

And so you do want to, if an SSD is offline, you want to store it in a cool place. But the lesson here really is, and this is the point that the article made. At one point it said don't assume that anything sitting on a shelf for years will be readable when you need it. You really do need to periodically plug it in, make sure it's still readable, and I would rewrite it, just to strengthen the bits.

MIKAH: I need to go back, though, because did I hear you right in that you said every six months you're having to replace a drive? Is that...

Steve: Yeah. So I have four times four, I have 16 spinning drives. And, yeah, I would say every six months or so a drive says, okay, I'm hurting. And so the RAID sends me email, and I go and swap it. Now, I have not yet run SpinRite on this drive. The last time that happened, there was actually nothing wrong with the drive. It just had that one little burst of trouble, SpinRite rewrote the sector, then the drive was fine. And I put it back in the RAID array, and it hasn't failed since. So, you know, it's a very touchy failure which is not a big problem. But I have had SSDs actually just completely die. So that's why I'm no longer thinking, oh, solid-state, that's way more reliable.

It's like, I guess what I'm saying is I refuse to have any loss of data. So refusing to have any loss of data means double redundancy and just keep swapping. So drives are consumables is the way I guess I would think about it, is they're consumables, and I'm consuming a drive at the rate of about maybe one or two a year in order to be running four RAIDs with four drives per RAID. And I've never lost a byte of data.

MIKAH: There you go. That's the brag that you get to at the end. Like, yes, I may be doing it every six months, but never lost a byte of data.

Steve: Okay. So this story is really fun. The Navy Times recently blew the lid off an intriguing story of a U.S. Navy warship. There were some officers who had installed a secret Starlink-based network on board so that a select few of the upper echelon would not be deprived of their precious Internet connectivity while they were deployed at sea. In their piece headlined "How Navy chiefs conspired to get themselves illegal warship WiFi," the Navy Times wrote the following. And I'm just going to share the beginning of it. It's a long article. I've got the link here for anyone who wants more. But here's what they said.

They said - this will give you the gist: "Today's Navy sailors are likely familiar with the jarring loss of Internet connectivity that can come with a ship's deployment. For a variety of reasons, including operational security, a crew's Internet access is regularly restricted while underway, to preserve bandwidth for the mission and to keep their ship safe from nefarious online attacks. But the senior enlisted leaders among the combat ship Manchester's gold crew knew no such privation last year, when they installed and secretly used their very own WiFi network during a deployment, according to a scathing internal investigation obtained by Navy Times.

"As the ship prepared for a West Pacific deployment in April 2023, the enlisted leader onboard conspired with the ship's chiefs to install the secret, unauthorized network aboard the ship, for use exclusively by them. So while rank-and-file sailors lived without the level of Internet connectivity they enjoyed ashore, the chiefs installed a Starlink satellite Internet dish on the top of the ship and used a WiFi network they dubbed 'Stinky' to check sports scores, text home, and stream movies."

MIKAH: Oh, come on.

Steve: "The enjoyment of those wireless creature comforts by enlisted leaders aboard the ship carried serious repercussions for the security of the ship and its crew. The investigation noted: 'The danger such systems pose to the crew, the ship, and the Navy cannot be understated.' Led by the senior enlisted leader of the ship's gold crew, then-Command Senior Chief Grisel Marrero, the effort roped in the entire chief's mess by the time it was uncovered a few months later.

"Marrero was relieved in late 2023" - and there was a court-martial, by the way - "after repeatedly misleading and lying to her ship's command about the WiFi network, and she was convicted at court-martial this spring in connection to the scheme. She was sentenced to a reduction in rank to E-7 after the trial and did not respond to requests for comment for this report. The Navy has yet to release the entirety of the Manchester investigation file to Navy Times, including supplemental enclosures. Such records generally include statements or interview transcripts with the accused.

"But records released so far show the probe, which wrapped in November, found that the entire chief's mess knew about the secret system, and those who didn't buy into it" - literally buy, b-u-y, into it - "were nonetheless culpable for not reporting the misconduct. Those chiefs and senior chiefs who used, paid for, helped hide, or knew about the system were given administrative nonjudicial punishment at commodore's mast, according to the investigation. All told, more than 15 Manchester chiefs were in cahoots with Marrero to purchase, install, and use the Starlink system aboard the ship.

"The investigation said: 'This agreement was a criminal conspiracy, supported by the overt act of bringing the purchased Starlink onboard USS Manchester. Any new member of the CPO Mess which then paid into the services joined that conspiracy following the system's operational status.'

"Records obtained by Navy Times via a Freedom of Information Act request reveal a months-long effort by Marrero to obtain, install, and then conceal the chiefs' WiFi network from superiors, including the covert installation of a Starlink satellite dish on the outside of the Manchester. When superiors became suspicious about the existence of the network and confronted her about it, Marrero failed to come clean on multiple occasions and provided falsified documents to further mislead Manchester's commanding officer, the investigation states. Unauthorized WiFi systems like the one Marrero set up are a massive no-no for deployed Navy ships, and Marrero's crime occurred as the ship was deployed to the West Pacific, where security concerns become even more paramount among heightened tensions with the Chinese.

"While Marrero claimed the WiFi system was secretly installed for morale purposes, the investigation notes that such a claim 'is undermined by the selective availability of the WiFi and strict control of its access to the CPO Mess only.'

"The Manchester's secret WiFi network was born in March of 2023, when Marrero and a co-conspirator got to work buying and installing the Starlink system before the ship's deployment began the following month. The Starlink dish was installed on the Manchester's O-5 level weather deck during a 'blanket' aloft period, which requires a sailor to hang high above or over the side of the ship. During a 'blanket' aloft, duties are not documented in the deck logs or the officer of the deck logs, according to the investigation. It's unclear who harnessed up and actually installed the system for Marrero due to redactions in the publicly released copy of the probe, but records show Marrero powered up the system the night before the ship got underway to the West Pacific waters of U.S. 7th Fleet.

"Marrero and her cohorts paid $2,800 for a Starlink High Performance Kit with a personal credit card, and contacted Starlink to expedite shipping so the system would arrive in time for the deployment."

MIKAH: I'm glad our tax dollars weren't used to purchase it. I'm glad to hear that.

Steve: At least it was paid for with personal money; right.

MIKAH: Yes. Whether it was then sort of, you know, came out in the wash, I don't know. But at least that [crosstalk] money.

Steve: "Starlink offers plans ranging from $90 to $5,000 a month, and allows users to control network settings via a cell phone app. The Navy is installing such authorized capabilities aboard some ships in the fleet. But that was not the case aboard the Manchester, where Marrero set up payment plans for the chief's mess to pay for the system either $62.50 a month or a one-time fee of $375."

MIKAH: It's a whole business. It's [crosstalk] WiFi.

Steve: "So that the ship's Chief Petty Officer Association treasurer collected the money into a chief's mess checking account."

MIKAH: What? This is a - this is wild.

Steve: "Those involved also used the Chief Petty Officer Association's debit card to pay off the $11,000 monthly Starlink bill." Ah, so they went for the $11,000 a month...

MIKAH: A thousand, $1,000, not $11,000.

Steve: Oh, sorry, sorry. Yeah, $1,000 a month Starlink bill. "And Marrero warned the chiefs to only use the network in their rooms. Marrero served as the gatekeeper of the system, records show, downloading and maintaining the Starlink app from her phone and naming it 'Stinky.'"

MIKAH: [Crosstalk] Kingpin.

Steve: That's right. "Only she could add others to the network and would directly type the password into their devices" - so that they would know what the password was. So it's like, don't worry, I'll type the password in, then you just use it, and it'll remember it. "After Manchester got underway from San Diego, Marrero and the chiefs soon realized the WiFi signal" - oh, darn - "didn't cover all areas of the ship."

MIKAH: Maybe because it's a lot of metal everywhere?

Steve: You think? "So the senior chief purchased signal repeaters and cable at the Navy Exchange store in Pearl Harbor, Hawaii, during a port visit in late April or early May, according to the investigation." That's right, we need to get us some repeaters here so that we'll be able to have access to Stinky no matter where we are.

MIKAH: Honestly, I'm becoming impressed. This is like multilevel - it's kind of getting impressive.

Steve: So they said: "Little stays secret within the close quarters of a deployed ship. And shortly after" - you can imagine; right? "And shortly after getting underway..."

MIKAH: Yeah. How are you on your phone right now? How are you watching that video?

Steve: Uh-huh. How do you know what the sports scores were?

MIKAH: Yeah, exactly.

Steve: "Scuttlebutt started swirling among some sailors about the unauthorized WiFi network. The ship's former executive officer, Commander Matthew Yokeley, caught wind of the rumors in May and notified the commanding officer, Commander Colleen Moore. Moore confronted Marrero about whether the chief's mess had an unauthorized WiFi network that same month. Another unidentified crew member approached Marrero about a WiFi network aboard the ship after finding available networks on a device that started with the name 'Stinky.' It's unclear who found the 'Stinky' network, due to redactions in the report.

"In both instances, Marrero denied that such a WiFi network existed. But she soon changed the 'Stinky' WiFi network name to another moniker that looked like a wireless printer, even though no such general-use wireless printers were present on the ship, the investigation found." Anyway, wow is right. Wow.

MIKAH: Wow. That's just very impressive. See, and here's the thing. I was very pleased to hear that our tax dollars were not used to make the purchase of this. But the money that went into the investigation of this and the subsequent court-martial and everything that was involved there, we did pay for that. So in that way I'm a little bit perturbed. Although, yeah, like, this is impressive. And I would love to hear Marrero's sort of reasoning, you know, because you've got to - is it like we just really wanted to watch our shows? Whatever it happened to be, you know, there's got to be more to it.

Steve: Right. You know, with so many people today seemingly unable to separate themselves from the Internet, it's foreseeable that there would be significant pressure to maintain connectivity while at sea. But at the same time we know how true it is that any form of Internet connection would need to be highly filtered. And that's the problem. You know, can you imagine

how much any hostile foreign power would love to get into a smartphone or a laptop of someone onboard who clicked the wrong link? You know, they're running TikTok, which is controlled by the Communist Chinese government indirectly. And then once in there, pivot with that access and jump into the ship's internal networks.

MIKAH: Yeah, well, and there's something to be said, too, for the individual who's behind the company that is responsible for Starlink in the first place also having that information of where that boat is at any given time. That's also something to consider. And it kind of makes you wonder if Marrero heard that these other ships were having this done and was like, well, I'm not waiting nine years to finally get it installed on this ship. We'll just do it ourselves.

Steve: So, but you bring up a good point that hadn't occurred to me, like the underlying motivation. Was it because she desperately needed to maintain access to friends and family or something?

MIKAH: Yeah, was someone ill or...

Steve: Or was it actually a profit center?

MIKAH: Right.

Steve: Was it like, hey, I can make some coin here by selling access to my fellow upper echelon.

MIKAH: Right. And I could also be simply, I mean, I can't speak to what all is involved, obviously, in any of these, what, patrolling bodies. But maybe, you know, if you are trying to recruit certain people, it's like, do we want to go to the place that has the WiFi, so whenever we go out on these long voyages? And it's like, look, we all have to put some money in because it's expensive. But, hey, at least you've got access to WiFi, and you don't have those pesky Navy blockers on everything, so you actually get to look at your TikTok. You know what I mean? Like there could be people who...

Steve: We heard that this is the party ship.

MIKAH: Yeah, this is the party ship.

Steve: So this is where we want to be.

MIKAH: You've got to pay a little bit, but it's still the party ship.

Steve: And I think, you know, another interesting question here is that it really, I mean, I get how addicted people - and that's the right word to use; right? - how addicted some, especially some people, not everybody, but are to 24/7 Internet access. And you're a sailor going out on a ship, and you're dark. You know, I mean, it's over. And the other thing is, how big and complex must the operation of these ships be that a Starlink antenna box sitting - it can be installed on the deck of the ship, pointing at the sky, and nobody walking by says "What is this box?"

MIKAH: What is that? I don't remember that being there. Like it might be a spy - like, come on. Yeah, those things have got a lot going on, apparently. Wow. And I've got to say, too, I sort of feel like Marrero could end up hosting one of our shows, given that Marrero knew - this is very clever, to change it to make it look like it was a wireless printer. That's clever. That takes knowledge of some level of networking, along with everything else that was involved, to think about being the one to do the WiFi pass, I mean, there's some intelligence there. Maybe eventually Marrero will make her way into, I don't know, the Pentagon or something.

Steve: I have a friend who named his own WiFi "NORAD scanner" or something. It was like, nobody's going to dare touch this.

MIKAH: Yeah, I'll leave that one alone, thank you. CIA watching you. No, I'll skip that one.

Steve: Okay. Next we're going to talk about the winner of a recent competition for the best

secure access edge service.

MIKAH: All right. We are back from the break, and let us continue on.

Steve: Okay. So it was, well, I thought it was a few months ago, turns out it was more like 10 months ago. I shared how impressed I was after meeting with and learning about the technology that the guys at ADAMnetworks, which is Adamnet.works, had created to help secure the Internet-facing border of enterprises. As I've noted several times, that is a daunting task, and it's not a job I would want. You know, like how do you secure Sony, you know, Entertainment when anybody clicking on a bad link can infect the network. So while I was delighted to see it, I was not surprised to discover that SC Magazine, a well-known and reputable security industry publisher, after running a head-to-head competitive comparison and evaluation of the industry's many various solutions, picked adam:ONE, which is the name of this system that they offer, as the winner.

SC Magazine wrote - and this is just a one-liner. They said at the top of their coverage of this, they said: "ADAMnetworks has claimed the prestigious Best SASE [S-A-S-E]" - I have what that acronym stands for here somewhere. It's not coming to mind. Oh, Secure Access Service Edge Solution. So "ADAMnetworks has claimed the prestigious Best SASE Solution award at the 2024 SC Awards for its cutting-edge product, adam:ONE. In a cybersecurity landscape where traditional reactive methods often fall short, adam:ONE stands out by providing a proactive, zero-trust security solution designed to eliminate threats before they infiltrate networks. This recognition places ADAMnetworks among the leading innovators in the increasingly competitive Secure Access Service Edge market."

So anyway, I just wanted to follow up on what we initially talked about back then, since protecting the enterprise from all of the mischief that those inside the enterprise might get up to is no small task. And since the adam:ONE folks appear to have the best handle on doing that job, I have a link to the SC Magazine's announcement with many more details in the show notes. And it was Episode 946, which was Halloween, October 31st last year, 2023, when I shared the results of my meeting with the adam:ONE folks. So anyway, for anyone who's interested, these guys, you know, won the best-in-class award for their solution for securing the enterprise perimeter. And I'm not surprised because, as I shared back at the end of October last year, they figured out how to do it, and they do it right.

Many of our listeners may have received email from me about the availability of SpinRite 6.1. Of course that won't come as news to anyone listening. But for anyone who purchased SpinRite 6.0, which was released 20 years ago, back in 2004, it would likely come as a welcome surprise. So over the weekend I received a note about this from a listener named Patrick. He wrote: "Good morning, sir. Quick note to let you know I've received an email from SpinRite.news." That was the domain that I used for sending the email. He said: "But it was flagged as spam by Exchange and dumped into my junk folder." He says: "Otherwise, thanks for the work on SpinRite 6.1. I'll let you get back to work on 7.0 now. Signed, Patrick."

So I replied to Patrick, writing: "Thanks for your feedback, Patrick. Since I'm mailing to all past SpinRite owners for the past 20 years, I'm sending those announcements through that domain you noted, SpinRite.news. Since that domain has not earned a reputation as a valid email sender, Apple is bouncing all incoming email addressed to anyone at me.com, icloud.com, and mac.com. And as you note, Exchange is routing incoming mail to spam," I said, "but at least it's not bouncing the mail back."

So I said: "My primary goal for this is twofold: I do want to inform any non-podcast listeners of the availability of a free upgrade to SpinRite 6.1, and I also want to remove all bad email addresses from the list for the future. This is why I'm sending from the previously unused domain SpinRite.news, since I expect that the bounce rate will be high, especially for the oldest 20-year-old email addresses, and I want to keep GRC.com's email reputation as spotless as possible." I said: "Once I've managed to update GRC's creaky old SpinRite owner list, I'll be able to mail from GRC.com, using its clean email reputation." I said: "Mail should then get into people's inboxes."

So it's been an interesting - it's been an education. I set up a relationship with the guys at Postmark that are an email forwarding service, and established all of the proper credentials and

security and cryptographic signing and everything to authenticate email coming from SpinRite.news. And then I began mailing from the most recent in the direction of the least recent. So from 2024, and then 2023, 2022, 2021, and so forth. And just as you'd expect, as I went back further and further in time, the bounce rate began to increase because people had, you know, left the companies where they were when they purchased SpinRite, so their inbox was terminated, and so forth. So I got as far back as 2011. I got all the way back through 2011, and the bounce rate had then reached one out of 10, which is the highest level that these guys are comfortable with me having, sending email through them.

So anyway, I'm going to come up with a different approach. The interesting thing, I mentioned that all of the email into Apple bounced. I saw the same thing occurring when I was first doing the regular podcast mailings like I did this morning. I sent out 9,400-plus pieces of email for this podcast to the subscribers to the Security Now! list. And for the first few weeks, Apple flatly blocked all of that. But then it stopped. And for the last handful of weeks, everything's been working perfectly, and email's been flowing without any problems.

So anyway, I thought it was just interesting. We've talked many times about how, when software is digitally signed, reputation is everything. Anybody can get a bad certificate, and bad guys do. You really have to sign your software now. But if it's not a certificate that has earned itself a reputation as signing non-malware, you know, Windows looks at it with a raised eyebrow and says, I don't know if I'm going to let this run and, you know, quarantines it. Turns out reputation is just as important for email. So my using - in retrospect, I'm not sure that I shouldn't just have sent from GRC because I'm able to see what the bounce rate is and throttle back or stop. But I didn't know how it was going to go. It actually kind of went better than I expected it to. You know, I was able to get all the way back through 2011. And now I need to trickle out emails the way I'm going to handle it, but from GRC, I think, and just make sure that our reputation stays good. But anyway, sort of an interesting experience.

The one thing that there is a kind of a call to action that I would ask our listeners, all of those who are listening who own SpinRite and, well, I guess have purchased it ever, but certainly since 2011, check your spam folders and see whether an announcement of SpinRite is in your spam folder. And, if so, mark it as not spam. I would appreciate that because that's the way we train the ISPs who are filtering that this is not spam I'm sending. It's, you know, I got your email address because you bought a copy of SpinRite once upon a time. So anyway, I would appreciate it if people check their spam folders and let their ISPs know, nope, that's not spam. I wish I had that in my inbox. And of course we'll be using those email addresses in the future for other good useful announcements.

And one last point, speaking of SpinRite. I have seen several people who have said, "Hey, you said that SpinRite 6.1 speeds up SSDs." One guy sent me the three benchmarks that he had from before running SpinRite. It was 131 MB/s. Then the middle of his drive was 184, and the end was 185. He ran SpinRite on it, and it came out at 120 MB for all three measures. In other words, from 131 down to 120 at the beginning, and then the ends was 184 and 185, also all now at 120.

So I replied to him to explain what happened. I wrote: "It appears that your SSD was mostly empty. So what happened is that those 'pre' SpinRite benchmark readings were illusory and were not really returning results from reading from the drive's physical media. SSDs and spinning SMR (Shingled Magnetic Recording) drives are aware of whether anything has ever been written to individual regions of their media. If nothing has ever been written, then there's nothing to be read. So they don't bother actually reading anything, since nothing is there other than blank space, you know, all zeroes or all ones, whatever they initialize to. So they just return zeroes or ones at lightning speed, at the full speed of the interface that connects the drive to its computer.

"But when SpinRite re-wrote the SSD's entire surface, the drive now believes that all of the media is now 'in use,' even though it may still only be storing all zeroes or ones. Now the drive believes that data is important to its owner. So when SpinRite's benchmark is run afterwards, what will be shown is the true reading speed from the media, which was exactly 120 MB/s everywhere.

"Now, after running SpinRite and remounting the SSD in an operating system, the OS itself will 're-trim' the SSD. It runs through the entire SSD's 'region in use' table, marking all of the regions that are not actually storing any file system data as 'empty' again. All current operating systems

do this periodically, so this will cure itself. Under Windows this happens weekly for SSDs, although you can run the Windows Disk Optimizer on the drive to cause Windows to do this on demand. And if you then re-run SpinRite's benchmark on the SSD, you'll see that the results have been returned to what they were before, so you can usefully then compare them with what SpinRite showed for its pre-benchmark results." So anyway, yeah, isn't that cool? So the drive's saying 184, 185, it wasn't actually reading anything. The drive actually reads at 120MB from the media. But when it knows that nothing was ever written in a region, why bother reading it?

MIKAH: Meaning, when you say it knows that nothing has ever been written, does that mean it has knowledge of the fact that it is a factory-fresh drive?

Steve: Yes.

MIKAH: Okay. And so does that suggest, then, that if I got a drive and then plugged it in, and it did not read at those much, not much higher, but higher speeds, then maybe that one was quality control tested or something at the factory? Or is there just kind of like a, right before it goes out into the world it gets a little blessing that says this is absolutely fresh, it's never been used?

Steve: It's exactly the second thing that you said. So there's a thing known as "trim." It's called "trimming the drive." The drive, with a resolution of some level of blocks, some number of sectors, it maintains a bit flag that knows whether any of the sectors have ever been written to that drive. And it turns out that it's really cool the way this works. The reason it's done is that the way flash memory, which is the NAND memory, the solid-state memory used in SSDs and in flash drives, you know, thumb drives works, is when you erase the data, all of the bits are set into one state. Like typically they're all set to ones. So in a completely empty region, they're all set to ones. But erasing could only be done in relatively large blocks, the same size as there is a bit for the trim of the block. So the whole block is written to all ones.

Now, say that you write a sector of data into that block. Writing the data sets a bunch of those ones to zeroes, but it's not possible to set the zeroes to ones. That is, only the process of erasing the entire block is able to set all of the bits to ones. So writing data is essentially like pushing the bits down. All that can be done is like to push them down to zeroes. And so if the system knows that regions have never been written to, then it knows there are no zeroes there, so it doesn't need to take the time to erase those. So it turns out it's a performance-enhancing and wear-minimizing sort of like background management of the data stored on the drive. And it's completely transparent to the user. You never see that happening. You know, you just store your data on the drive, and everything works.

But in the background, Windows is making sure that, like, so you delete a file from the file system. And, you know, back in the day we know that you could undelete things. That was what Peter Norton discovered that made him famous because he came out with Unerase or Undelete where you could just, you know, get your file back. But today's file systems actually release the space. So they release the space. But the space, as we know, the file is still there until it's actually like physically rewritten. Or until Windows says to the SSD on Sunday night, or whenever the so-called "optimization" occurs, Windows will say, oh, by the way, that chunk of region no longer contains any actual data. It's no longer part of the file system. And so then the drive takes advantage of that information to know whether or not it needs to preserve it when it's doing its normal work. So a lot going on behind the scenes that people don't normally take a look at.

MIKAH: Yeah, that is fascinating to me, honestly. I didn't realize that it was kind of just, yeah, as you said, sort of you don't to worry about it, but I'm going to worry about it and will make sure that we're not doing more than needs to be done. If it ain't broke, don't fix it sort of situation.

Steve: Right. Yup. Well, and my favorite reaction to the weekend's mailing was just a one-liner from someone named Jim. He said: "Wow." He got the email. He said: "Talk about turning back time to the days of Madonna and a portable CD player that took 'D' cells." And he said: "Love it." He says: "I bet if I dig long enough, I might be able to find a copy on floppy of SpinRite from early 1990 or so." And he said: "Awesome. Thanks. You saved my tail many times in the past."

MIKAH: Aw.

Steve: So anyway, very cool. And finally, one piece of Closing the Loop feedback. I mentioned this to you before the podcast, Mikah. Chris Paetz wrote: "Just listened to your podcast and heard you say that you won't be able to find your email address anywhere online. But Perplexity.ai knows where to find it, apparently, FYI." He said: "I asked the question 'What is Steve Gibson's email address at Security Now!' and received this reply: 'Steve Gibson's email address for the Security Now! podcast is securitynow@grc.com. This address is intended for podcast feedback and is mentioned regularly during the show, although it is not prominently displayed on the Gibson Research Corporation website.'"

MIKAH: Whoa.

Steve: And I'll just say that I continue to be surprised by what we're creating here with this what can only be considered an AI revolution. As we humans continue to push technology further and further, we create both new capabilities and new dilemmas. You know, we're rapidly moving into a world that was only recently pure science fiction. So it's going to be truly interesting to see what happens with AI. But, Mikah, like, these sorts of summaries, where, like, it appears to actually get the context and, like, to understand what's going on, it's creepy.

MIKAH: Yes. It gets a little creepy. And honestly, for me, I remember when we first started seeing the generative AI thing take off. And I was still in a head space of, I know it's kind of silly, but I would feel bad about challenging the system too much. I would feel as if I was, you know, making it - I didn't want to give it the opportunity to fail me because I don't like to - this is kind of - but I don't like it when I can't, you know, properly achieve a task. And so I would be overly descriptive and overly mindful and try to help it get to the right answer that I was looking for because I didn't want to be disappointed, but I also didn't want it to, like, work hard. I don't know what was going on psychologically, but that's just where I was. And now I'm very lazy about how I'm asking the things because it does, you know, on the times whenever I'm asking it to...

Steve: It appears to understand.

MIKAH: It appears to understand. And I think that that's really going to make all the difference when we see more of this playing out in our virtual assistants. The actor Bella Ramsey has just - Apple just debuted several advertisements featuring Bella Ramsey. And they are an actor in "The Last of Us." And they were showing off some of the aspects of a new version of Siri where you could say "Where do I know this person from," and it was able to say, oh, yeah, you met them because I can look at your calendar, I can look at this. And up to this point, you know, we haven't - you have to do so much of the directing for that to happen.

Steve: Yes, yes.

MIKAH: And so it is really wild seeing how this is just happening automatically in the background. And then when you use something like Perplexity, because what Perplexity does is it is a search engine that has AI tacked onto it. And so earlier when we were talking about it, Benito had pointed out that that was probably a big part of that is it has access to the transcripts for Security Now!, which that alone is something that needs to be, you know, if it has access to those transcripts because it does a search on finds the show on the TWiT website and can read through that, it has information that it wouldn't otherwise have.

And another quick thing that I'll say, and then I'll get off this topic, is I was just - you may have seen that Google had updated its NotebookLM system. NotebookLM is the online tool where you can give it your Google docs or your other Google files. And then you can say, for example, I would, for example, say "What were the articles that I wrote back in 2015 where I talk about this Smart Home product," because I was doing that at the time.

Steve: Wow.

MIKAH: It looked back at my documents and helped me find those things. And then I could say "Help me break down a summary of what my thoughts were at the time." Right? But here's where it gets wild, Steve, is just recently NotebookLM released a new sort of feature of this where you feed it your documents, and it makes a podcast conversation of the documents that you give it.

So it has two people that are voices that are completely AI-generated, having a conversation about specific documents.

And there was an app developer on, I want to say it was Mastodon. And I'm sorry I don't have the links to this right off the top because I didn't know I was going to be talking about this. But basically this developer wanted to see what would happen if they fed it some weird stuff and so just gave it, like, 35 documents, or maybe even was more than that. They were all just named Patent One, Patent Two, Patent Three, Patent Four, and just had binary in it, all zeroes and ones. And there was a somewhat compelling fake conversation taking place between these two AI voices about these Patents documents that weren't real. Mind-blowing stuff.

Steve: So, like, the perfect example of an AI hallucination.

MIKAH: Yes, yes.

Steve: Just completely made up out of whole cloth.

MIKAH: Yes.

Steve: This, you know, stuff that sounded absolutely convincing.

MIKAH: Because there were these little interruptions like we're doing here, where I'm kind of going "yes" and, you know, making those sounds. It was doing that, too. And so it was chilling. It was honestly chilling. But at the same time it was hilarious because it felt like an SNL parody almost. And it was - I was feeling a lot of emotions all at once. So I'm sure that you were feeling a little bit of that when this came up because it's like, what are we building here?

Steve: Well, and what you just described about Google LM and the idea that it can look at your document archive, your historical document archive, and be, like, bring search to, I mean, an entirely new meaning to the term "search," this is why Microsoft is so desperate to get Recall into everyone's Windows machine. They want Windows to be taking a snapshot of everyone's screens every couple seconds, understanding and saving, archiving what's there, for exactly the same purpose.

So future Windows users who do not object to the concept, I mean like the privacy implications of having their machine and Microsoft to some degree having access to the entire, you know, like unredacted contents of all the screens that have been on their machine for years, they're going to be able to pose these sorts of questions. I mean, so it would be not only Google Docs because those Google Docs were on your screen at some point, and Recall would have taken pictures of them, but everything else that you did with your computer. And, you know, you want it to be highly confidential, to be just, I mean, just between you and your computer. But one can imagine the power that it would have. And it is, I think, terrifying.

MIKAH: Can you imagine credit card companies getting some sort of proprietary score from your behavior that is generated by - and I know I'm going a little out there. But at the same time, we've seen with health insurance companies and auto insurance companies who have those little dongles, or you get an Apple watch from them, and then you have to send it your stats, and then over time it can reduce your bill, imagine that same thing applied where your social credit is generated based on how your - think of it like, can you imagine a corporation that has cybersecurity insurance, and it uses a score generated by Recall to see how risky your employees' behavior is on all of your company-owned devices.

Steve: Yup.

MIKAH: Right there, that can mean, you know, 10,000 fewer dollars paid every year because your employees are properly trained, and the behavior plays out. And then, you know, the in-between, the company that's doing the Recall part can say, we're not going to give the company exact information. We've generated a proprietary score that says there are 98 out of 100 mice icons, you know what I mean, whatever the score happens to be, that's interesting stuff.

Steve: Well, and we've seen and covered on this podcast instances where cars are now feeding back their drivers' driving habit data.

MIKAH: That's right, yes.

Steve: And it's affecting their insurance rates.

MIKAH: Yup, yup, I remember you talking about that.

Steve: So this is not farfetched.

MIKAH: No.

Steve: We know that there will be tremendous pressure to obtain this information. And Microsoft is all about generating revenue.

MIKAH: It's a company, folks.

Steve: Uh-huh.

MIKAH: It's all about the money now.

Steve: Yeah, we are the product.

MIKAH: Yes. Yes, indeed. Yes, indeed.

Steve: Okay. So today's exploration topic began with my receipt a week ago of a note which read: "Hi, Steve. My name is Ben. I am a former Ben-Gurion University student and long-time listener of Security Now!. You covered a few studies of mine in the past including: Lamphone (speech recovery from light bulbs), Video-based Cryptanalysis (key recovery from a power LED), and Morris-II (the AI worm).

"In two recent studies," he says, "that I co-authored and were led by Andres from Cornell Tech, we revealed attacks against end-to-end encrypted applications, demonstrating the recovery of encrypted confidential data from backups of two messaging applications (WhatsApp and Signal), and 10 password managers (LastPass, Dashlane, Zoho Vault, 1Password, Enpass, Roboform, Keeper, NordPass, Proton Pass, and KeePassXC). We named these attacks 'injection attacks,' and the papers were published on USENIX Sec '24 and Security and Privacy '24. Attached are the links I believe that your audience will find interesting, as once again they prove that while end-to-end encryption is the best approach for applications, the devil is in the implementation."

Now, of course, I had seen from Ben's note envelope that his full name was Ben Nassi, which I recognized immediately since, as he noted, we've covered all of his work and exploitation discoveries through the years. So needless to say, I was interested in what new mischief Ben and his fellow security researchers had gotten themselves up to now.

Five or six researchers, depending upon which paper, led by Cornell University's Andres Fabrega, from Cornell University and Cornell Tech, collaborated on two papers. One was titled "Injection Attacks Against End-to-End Encrypted Applications," and the other was titled "Exploiting Leakage in Password Managers via Injection Attacks." To get us started on our understanding of what they have done, I'm going to share the Abstract from the paper about Password Manager injection attacks.

It's pretty brief, and it says: "This work explores injection attacks against password managers. In this setting, the adversary only controls their own application client, which they use to 'inject' chosen payloads to a victim's client via, for example, sharing credentials with them. The injections are interleaved with adversarial observations of some form of protected state, such as encrypted vault exports or the network traffic received by the application servers. From this, the adversary is able to obtain confidential information.

"We uncover a series of general design patterns in popular password managers that lead to vulnerabilities allowing an adversary to efficiently recover passwords, URLs, usernames, and attachments. We develop general attack templates to exploit these design patterns and experimentally showcase their practical efficacy via analysis of 10 distinct password manager applications. We disclosed our findings to these vendors, many of which deployed mitigations."

Okay. So that's interesting. When they use the term "injection," they're referring to providing an unwitting target some information that the target will cause to be stored in their own instance of their password manager. And then by examining the target's encrypted vault or their network traffic, presumably synchronizing their encrypted vault with the password manager's cloud backup, they're able to learn as much as the user's secret username, passwords, URLs, and attachments. So we need to learn more about that.

But I'll first note that the second paper does something somewhat similar with end-to-end messaging, specifically WhatsApp and Signal. In the case of messaging, an attacker sends messages to a targeted user. Assuming that the attacker is somehow able to obtain and observe the targeted user's encrypted cloud backup, the researchers were able to demonstrate their ability to determine whether the target had received specific attachments; or, for example, which of two messages the target had previously received.

Now, clever as these researchers are, I came away feeling, I guess, better rather than worse about the safety of WhatsApp and Signal. My feeling was like, okay, wow. If that's the most intrusion that these guys were able to achieve given their obviously serious skills, then that says a lot about how good these apps are. But that said, since the goal is true zero leakage of any kind, I would not be surprised to learn that the app vendors had added something like length fuzzing to the things being stored, specifically to thwart the leakage that these guys were able to induce. So the research was definitely useful.

MIKAH: So in other words, what you're saying is because they did this, even though it didn't give much information, it was still worth doing because it resulted in the vendors making changes to even make this small vulnerability even less likely to be effective.

Steve: Exactly.

MIKAH: Okay.

Steve: Yeah, certainly the goal is zero leakage. And we often quote Ben Schneier - not Ben. I can't believe I've forgotten his name, his first name. Bruce. Bruce Schneier.

MIKAH: Bruce Schneier, oh.

Steve: Yeah. Bruce said, and I love this, attacks never get worse, they only ever get better. Meaning, you know, I mean, it's obvious in retrospect how attackers are getting more clever, they're never getting dumber or less clever. So the attacks only ever improve. So we start with, like, okay, this doesn't seem very bad. But still.

MIKAH: The step one, yeah.

Steve: Exactly. It's like the first step in a new vulnerability is that it crashes the system. The next step is it no longer crashes it, you take it over.

MIKAH: Yeah.

Steve: So, okay. But the password manager leakages, from what we've seen so far, appear to be somewhat more dire. So let's take a closer look at them. These researchers set the stage in their introduction of that paper by writing: "Password-based authentication suffers from well-known pitfalls" - that is, just password-based, you know, password authentication, using username and password on a website.

They said: "...such as the fact that users tend to choose passwords that can be easily guessed by

attackers. Password managers are often cited as the default solution to this problem, as users can offload to them the complexities" - thank god - "of password generation, storage, and retrieval. Indeed, password managers have enjoyed a notable rise in popularity, placing them among the most ubiquitous of security-oriented tools." I can't believe, I can't imagine that anybody listening to this podcast is not an avid password manager user at this point.

MIKAH: Right.

Steve: So they wrote: "Password managers have benefited from academic attention, which has helped understand and improve their security along various dimensions. The attacks uncovered by prior work broadly fall under two general threat models. First are attacks that use a client-side resource controlled by the adversary, such as a malicious website visited by the client, a rogue application in the victim's device, or the client's own WiFi network. Second are adversaries that somehow acquire a copy of a user's encrypted vault, and exploit leakage from unencrypted vault metadata or by offline cracking attacks of a user's master password. State-of-the-art password managers are therefore designed to resist both kinds of threats, and notably use slow cryptographic hashing to prevent cracking attacks for well-chosen master passwords.

"In this work, we consider a new kind of threat model in which an adversary, one, controls their own application client, meaning their own password manager, through which they can send chosen payloads to the victim, for example, via the password sharing feature now found in most modern password managers; and, two, can observe some form of encrypted state and associated metadata, such as the user's encrypted vault backups or network requests received by the application servers. Borrowing terminology from prior work in other domains, we refer to attacks in this threat model as 'injection attacks.'

"The core idea behind injection attacks," they wrote, "is that the adversary can use injections to trigger subtle interactions in the application logic between their data and target victim data, for example, other passwords used by the target, which are reflected in their observations of ciphertext, for example, inspecting their lengths, and metadata in a way that allows recovered sensitive information. We argue that this threat model is increasingly important as password managers become more complex and feature-rich" - in other words, you know, more ways for things to go wrong because it's like, oh, let's add this. What could possibly go wrong? They said: "...which provides new avenues for injection mechanisms and vulnerable cross-user interactions.

"To understand whether this threat model is of practical concern or not, we performed a security analysis of 10 popular password managers that support sharing: LastPass, Dashlane, Zoho Vault, 1Password, Enpass, Roboform, Keeper, NordPass, Proton Pass, and KeePassXC. Together, these reportedly account for over 30% of all password manager users." And, okay, that's shocked me. That seems low.

MIKAH: I thought so, too, yeah. Where are the rest of them?

Steve: Yeah, I mean, LastPass, though they had that trouble, I thought they were a huge percentage of the total password manager install base.

MIKAH: Maybe they're Chrome? Where everybody's got their passwords is Chrome?

Steve: Got me. But that 30%...

MIKAH: And Apple Keychain I guess would be the most commonplace.

Steve: Yeah. So they said: "We uncover a series of exploitable vulnerabilities that implicate all of the password managers investigated. Our first class of attacks exploits the fact that a common feature of password managers is for clients to periodically log outside the device various metrics about the 'health' of a user's vault, such as the number of duplicate passwords. We show..."

MIKAH: Oh, good.

Steve: What?

MIKAH: Those are outside of the vault? That's great. Love to hear that.

Steve: Uh-huh, yeah. And so there's an example, right, of like what, you know, we're obviously done with the problem of storing usernames and passwords. What new features can we add? How can we further enhance this? And that's where we start getting into trouble, as we'll see. So they explain: "We show how an adversary can leverage these benign-looking metrics to perform an efficient binary search-based dictionary attack that recovers the target's saved passwords. Our attacks do not require the adversary to know additional information about the victim's saved credentials beforehand, for example, URLs nor usernames.

"Five out of the 10 applications are vulnerable to this attack. In most cases, the adversary must be a passive eavesdropper that observes these metrics directly, for example, by having a persistent foothold in the application servers; while for one application the attack is feasible by a passive network adversary that simply observes the HTTPS channels under which the end-to-end encrypted data is transmitted. We note that both eavesdropping and network adversaries are within scope for the threat models under which password managers are designed, meaning they weren't supposed to be leaking this information. And the ubiquity of server-side breaches, combined with the difficulty of detecting such breaches, make it critical that password managers resist, that is, are resistant to such attacks." So we'll come back to that in a minute.

They said: "Our second class of attacks exploits another feature of password managers. Clients often display a small identifying icon, such as a company logo, alongside each of a user's saved credentials. Importantly, such icons are only fetched once per URL, and subsequent credentials reuse the icon stored in the client. We show how this fact allows an adversary to perform an efficient dictionary attack on the URLs in a victim's vault. The attack always succeeds in our experiments, and mounting it requires no additional assumptions about the victim's saved credentials. Six of our 10 case study applications are vulnerable to this attack, and in all cases exploitation only requires observations by a network adversary," meaning someone able to watch the user's network traffic.

And then, finally: "We turn our attention," they said, "to adversaries that have an encrypted copy of the entire vault, such as compromising a local password-protected database file or backup of it. In this case, we analyze the security of KDBX, which is a file format used by many password managers, notably KeePass and its derivatives. To optimize for storage, KDBX employs a variety of storage-saving techniques, such as file deduplication and compression.

"We show two attacks exploiting these features to recover URLs, usernames, and attachment contents. Compression and deduplication have led to attacks against other systems before, but our work is the first to show that these types of vulnerabilities also arise in the context of password managers. Our attacks target features of the underlying file format itself, and thus can potentially be leveraged against any application that uses KDBX. We implement a proof-of-concept for our attacks in the case of KeePassXC, and experimentally show that its accuracy is sufficiently high to make it a practical threat.

"A summary of our attacks follows. They exploit common design patterns found in password managers. And as such, other applications that employ these can be vulnerable to our attacks. Indeed, for each of our attacks, we describe a general template for it, which is agnostic to lower-level application details, and that can be used to target any application that follows the relevant design pattern. More broadly, our findings uncover higher level issues in password manager design, and we discuss the future work that will be required to provide generally applicable mitigations for injection attacks."

Okay. So I want to clarify the nature of these attacks so that the vulnerabilities they found will make more sense. They identified three ways of attacking the security of today's password managers. The first class of attacks, which the researchers refer to as "Vault-Health Logging," rely upon the newer features of application-wide metrics. These arise from the behavior exhibited by many password managers which compute various metrics, like the example they gave, the number of passwords their user has duplicated in their vault, which includes both personal and shared vault entries.

MIKAH: Would this also include, Steve, the Have I Been Pwned integration that a lot of these password apps have that tell you, oh, yes, your password is part of a vulnerability. That's another health...

Steve: Yes, that's another perfect ability of like a new feature that can be, can inadvertently leak some metadata about what you've got in your vault. When these metrics are logged outside the device, such as by the application's cloud servers, an adversary - and that means that logging is visible on your network, even if it's encrypted, it still, by the size of it, is visible, then an adversary can induce fluctuations in these metrics with so-called "injections," and observe how they are updated in the external location. So to carry out these attacks, in addition to having the ability to inject credentials, this attack requires the adversary to have access to the location where the metrics are logged. Okay. So, you know, some sort of foothold of some kind is required.

So specifically, LastPass, Dashlane, Zoho Vault, Keeper, and NordPass are those five they mentioned before. Those are all vulnerable to these vault-health logging attacks because those five offer these features. All of the password managers except Zoho Vault require that server-side foothold be present. So, and that's also effective against Zoho Vault, but it's additionally vulnerable, Zoho Vault is additionally vulnerable to simpler passive network attacks. Okay. So the vault-health attack can be used successfully against Zoho Vault only when passive networking, or passive network eavesdropping is present.

The second class of attacks is, I guess I would call it, and they refer to it as "URL icon fetching." Actually, what's interesting, too, we've actually seen this in browser-based attacks where people's web browsers were caching the favicons of the sites they were visiting. So this is sort of like that. It arises from the fact that many password managers, as they said, display a graphical icon next to each credential, identifying the website associated with it. This icon will only be fetched once from the application servers, and future entries for the same website reuse the image from a client-side cache.

An adversary could use this to determine whether or not a target's password manager has previously obtained and cached a credential for that particular URL. If the attacker induces the target's password manager to get an icon, that means there was no previous entry for that website. In addition to having a means of updating the target's credentials, probably without credential sharing, this attack requires the attackers to be able to observe the HTTPS request traffic that leaves the victim's client, or to have a foothold in the server from which the icons are fetched. So the researchers found that Dashlane, 1Password, Enpass, Roboform, Proton Pass, and NordPass were all vulnerable to disclosing their protected URLs through this passive network eavesdropping.

And the third and final class of attacks only affects KeePassXC, that is, among the top 10 that were tested. This arises from KeePassXC's storage file system. KeePassXC uses this system to decrease the size of its encrypted vault. Now, the problem is the presence of data compression - and this is something we've talked about before, in secure storage, which is used obviously to reduce the storage redundancy - can be used to reveal the data that's already been stored in the compressed store. The size of the storage will increase when unique new information is stored. But if something is added that already exists in the storage, the compression that's present will keep the total storage from increasing.

MIKAH: Oh. You can use that, I mean...

Steve: And so you can infer, exactly.

MIKAH: Wow.

Steve: You're able to infer what was there based on whether it now occupies as much space, additional space, as the new information would have required.

MIKAH: Aha.

Steve: So isn't that a cool side-channel?

MIKAH: That is so cool. And so clever. These clever, clever people.

Steve: Yeah. So, you know, it's a side-channel that we've talked about in the past. It's necessary to be very careful about compressing secret data before it's encrypted because this side-channel can be used to leak information about the content of the secret data by inference. And we know that post-encryption compression is never used because it's not possible to compress encrypted information since anything that's been properly encrypted is indistinguishable from completely random noise, and completely random noise is, by definition, incompressible.

MIKAH: You know, I needed that, Steve, because I always - the logic to me there never quite lined up in terms of how in the world can I take a thing that exists in this digital space, like why is it not already as small as it can possibly be? Why does compression work at all? So the way that you've just described it there helps me to get why compression even works in the first place because, if you do see those patterns, then you can use those patterns to help make it smaller.

Steve: Yes.

MIKAH: I just thought, yeah, why isn't the JPEG even smaller than - why can I take some JPEGs, put them into a ZIP, and the ZIP is smaller? And now I understand.

Steve: Right. And the ZIP, we talked about this on the podcast, like, forever ago, two researchers who were at IBM, Lempel and Ziv, created something, and their initials are LZ, LZ compression, and that's the compression used in ZIP and a lot of the early compressors. Basically, as you're feeding data in, the compressor is storing the history of the data that is seen. And as new data comes in, it looks in the history buffer for any previous occurrence. So like even, like the word "the," when it sees "the," like it looks in this buffer of past stuff that it has seen uncompressed, and it says, oh, look, that was here. So instead of putting "the" in, it puts a pointer to where it is in the buffer because a pointer is much shorter than the word.

MIKAH: Got it, yeah.

Steve: So normally when you're talking about something, the context of your conversation is reusing the same words over and over within a short period of time.

MIKAH: Ditto, ditto, ditto, ditto.

Steve: Exactly, ditto, ditto, ditto, ditto. And so it just aligns all that. And all it does is put little pointers to where it recently occurred.

MIKAH: That's cool.

Steve: Instead of having to restore it. It is super clever. They obtained a patent at the time, but patents only last 17 years. So it's, you know, everyone's been able to use it afterwards.

Okay. So finally, the last of the attacks, which is this compression attack, KeePassXC is vulnerable to this type of attack because it does exactly that. By examining fluctuations in the size of KeePassXC's encrypted vault, after injecting known information, it's possible for an attacker to glean information about the vault's unknown contents. To pull this off, in addition to injecting credentials, this attack requires the adversary to have persistent access to the victim's encrypted vault, either directly or by monitoring vault backups.

In their work, the researchers demonstrated attacks against two mechanisms, both the compression and attachment deduplication. So if you gave it an attachment, you would see, if it expanded by the size of the attachment, you knew that it didn't already exist. If it didn't expand because the vault deduplicates, it was like, aha, that already exists somewhere. Again, doesn't seem like a big problem. But it is. You are leaking some information.

MIKAH: Right.

Steve: So after putting these 10 password managers through the wringer, they contacted each of the password managers' vendors to share their results ahead of taking this work public. Here's what the researchers reported from that effort. They wrote: "We reported our findings to the 10 vendors affected by our work, many of which proceeded to deploy mitigations. LastPass adopted our suggested mitigation of separating vault-health metrics between personal and shared credentials, which disables the injection channel. They released an initial implementation of this fix in version 4.129.0, removing shared folders from the vault-health logs, as these lead to the most severe variant of our attack. Removing individually shared credentials from the logs is more technically challenging - and individual credentials lead to a less practical version of our attack - and thus has been deferred to later in their roadmap. Their projection is to release this fix by the end of the year, which would complete a full mitigation of our attack.

"Zoho Vault plans to adopt a similar fix by implementing an option to separately compute vault-health metrics on personal passwords as of version 4.0. Dashlane opted for a partial mitigation instead, namely, increasing their rate limits on the sharing endpoints 'as much as possible.'" Okay, whatever that means.

MIKAH: Yeah.

Steve: "Given the fact that their vault-health metrics are only logged once per day, their tight sharing limits significantly affect the practicality and runtime of the attack." In other words, they're updating so seldom already that it would really slow things down.

MIKAH: You'd have to really be sitting around, yeah, watching constantly.

Steve: Yeah, right. "In addition, the resource limits on their web application and extensions prevent an adversary from sharing an unlimited number of credentials with a victim, which increases the runtime of the attack even more. As part of their disclosure, they informed us that incorporating shared passwords is a core feature of their vault-health metrics, and thus removing shared passwords would represent a notable disruption to this feature." So they don't want to stop doing that. Whereas LastPass said, yeah, we just took it out.

"Then, to address our URL icon fetching attack, Dashlane implemented a new feature as of version 6.2415 that allows users to turn off fetching credential icons, which disables the side-channel for both an eavesdropping and network adversary, and thus provides a full mitigation to our attack. To address our attack even when URL icons are turned on, they additionally migrated their icon fetching tool to a new endpoint, api.dashlane.com, which is used by multiple parts of their application logic. As such, this would make it significantly more challenging for a network adversary to use traffic analysis techniques to identify whether an icon fetch request is included in the traffic sent to this top-level endpoint, due to the high amount of noise from other requests sent to the endpoint."

In other words, it used to be that the URLs were being pulled from their own domain. And so they said, oh, no problem, we'll just add them to an existing domain which is under heavy use already. So the actual URL icon request will be lost in the noise. You won't, you know, an attacker who can just see the traffic happening without knowing what it is, because that is after all HTTPS encrypted, they just won't be able to determine. They'll just give up and go away, even though it wasn't that bad an attack anyway.

They wrote: "NordPass also adopted our suggested mitigation of separating vault-health metrics between personal and shared credentials, which disables the injection channel and is thus a full mitigation to our attack. This fix is planned to be deployed by the end of August 2024." And I guess that's already happened, then. "Then, to address our URL icon fetching attack, NordPass added a feature to disable URL icons by default" - by default, thank you - "which provides a mechanism to disable the injection side-channel. This fix is planned to be deployed before the end of 2024. In addition, they're currently exploring more robust mitigations for this attack, to protect users even when URL icons are turned on.

"To address our attack on attachment deduplication, KeePassXC adopted our suggested mitigation of deduplicating files separately for every shared folder, which disables the injection side-channel.

Then, to address our compression-based attacks, they modified their file format by, every time the database is saved, picking a random length between 64 and 512 bytes, generating a random array of this length, and including this in a 'custom data' field for their file format. We note that this is only a partial mitigation." Basically what that is is that's fuzzing. It's fuzzing the length so that there is no direct correlation between lengths any longer.

So they said: "We note that this is only a partial mitigation, as an adversary can potentially use statistical techniques to bypass the noise." Eh, okay, but not really. "This, however, would require a significantly higher number of injections. Both fixes were promptly implemented by the KeePassXC team, and have since rolled out as part of version 2.7.

"Enpass already provides support for turning off URL icons, which is off by default" - thank you - "and thus users who disable URL icons are not vulnerable to our attack. As a first step towards more mitigations, Enpass added an option for organization admins to control this setting via an organization-level policy. They've decided not to deploy mitigations at this time to address the attack even when URL icons are turned on.

"1Password" - a sponsor of the show as noted earlier - "already provides support for turning off URL icons."

MIKAH: Which I'm doing right now.

Steve: Uh-huh.

MIKAH: I did not know this. I'm turning it off.

Steve: "And thus users who disable URL icons are not vulnerable to our attack. However, 1Password decided not to deploy additional mitigations to address the attack at this time, even when URL icons are turned on. Similarly, Proton Pass already provides support for turning off URL icons, and thus users who disable URL icons are not vulnerable to our attack. We shared suggestions for how to address this attack even when URL icons are turned on, but do not have details on their plans to deploy these.

"Lastly, Keeper considers our attack on their system a very low severity issue" - and I can't really argue with that - "and opted not to deploy mitigations, and instead have added changing this feature as a consideration for an upcoming platform update. As part of the communication, they shared that removing transferred credentials from the count of duplicate passwords displayed in the Admin Console would represent a notable disruption to a feature of the business product."

Okay. So where does this leave us? I agree with Ben Nassi that this is an interesting and potentially important work. But it's clearly of more interest theoretically than practically. As a true threat, it's really out there on the fringe. It's not impossible that some extremely motivated attacker might somehow arrange to set themselves up in a position to pull off one of these attacks against a specific high-value target, but I would say it is safe to say that none of us listening to this have anything to worry about, even before these obscure holes were plugged.

So the reason I chose to share these attacks on this podcast is for what we learn from them about the true challenges that are associated with truly protecting secret information. It's so easy for the salesman to boast, "Oh, don't worry, it's all military-grade encrypted. And guess what! We're using a bazillion-bit key! So no one will ever possibly crack that." Right.

But the lesson taught by these injection attacks is that no one ever needs to crack that bazillion-bit key. The reason the password managers jumped to modify and improve their systems when they were informed of these subtle issues is that "subtle issues" may be all that's needed to infer the data that's being protected by those bazillion-bit keys. Any mature and fully informed understanding needs to appreciate that encrypting something is far from being the end of the task. The encryption is only the start. Twenty years ago there was a general lack of understanding of the gulf that exists between theoretical and practical "field-ready" security technology. But during these past 20 years this sort of research, exactly like we've just seen, has opened the eyes of people who are implementing these systems, and we have all benefitted.

So for anyone who may be interested in digging deeper, I've included the links to both of these research papers at the end of the show notes. And that's it.

MIKAH: One question I have for you because, if I don't ask it, the chat room will riot. Do you feel there's an omission that Bitwarden, also a sponsor on the network in the past, was not included as part of these? Or do you think that Bitwarden was not affected? Or what do you think about Bitwarden being...

Steve: That's a great question. I absolutely wonder that, why they didn't include Bitwarden in their 10. I got the sense that, like from reading between the lines, that some of these are more enterprise-oriented solutions.

MIKAH: Understood.

Steve: And so, not that Bitwarden can't be used and shouldn't be used in an enterprise, but that may have skewed their choice for some reason. But it's a good question. And what anyone could infer is that, if Bitwarden is being used to, like, pull icons from URLs, we don't know that that's a vulnerability, but it is a side-channel that all of these password managers have turned off or don't have turned on by default. So it's really a good question. I have no idea why Bitwarden was not there.