Windows Sandbox

Description: Enabling Firefox's Tab Grouping. Recalled Recall Re-Rolls out. The crucial CVE program nearly died. It's been given new life. China confesses to hacking the U.S., blaming our stance on Taiwan. CISA says what Oracle still refuses to. Brute force attacks on the (rapid) rise. An AI/ML Python package rates a 9.8 again. The CA/Browser Forum passed short-life certs. A wonderful crosswalk hack hits Silicon Valley. Android to add force restarting ahead of schedule. Maybe. The EFF is never happy, but especially now, about Florida. Interesting research into ransomware payouts. Windows Sandbox: The amazing gem hidden inside all Windows 10 & 11!

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-1022.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-1022-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. He's figured out how to enable Firefox tab grouping. He'll share that with you. Good news, Mitre's CVE program is not dead yet, and there are plans to keep it alive forever. And we'll find out about a Windows feature that's been there for a long time, but Steve has just rediscovered it. It's the hidden gem inside all versions of Windows 10 and 11: The Windows Sandbox. That and, yes, it's time for short-lived certs. All that and more coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 1022, recorded Tuesday, April 22nd, 2025: The Windows Sandbox.

It's time for Security Now!, the show where we talk about your privacy, your security, staying safe online, all of that courtesy of this cat, Mr. Steve Gibson, "Live Long and Prosper" Gibson.

Steve Gibson: That's right.

Leo: At GRC.com. Hi, Steve.

Steve: I can do it with both hands.

Leo: I can't. I can't. I have to have my fingers taped. Actually, you know what?

Steve: And that's a thing, isn't it.

Leo: Now that I've been playing the piano.

Steve: Oh. Because they're limber digits.

Leo: Well, it's one of the things you have to learn with the piano because you've got tendons connecting your pinkie and your ring finger, and you've got to isolate those. You've got to learn to isolate those. So maybe I'm a little better at the...

Steve: You're the pinkie isolator.

Leo: I am the - you do these exercises, and then you do the Hanon things, and there's all these things you have to do. Hey, before you get to what's coming up on today's show, I want to show you.

Steve: They say it's good to learn new things.

Leo: I'm trying to keep my brain...

Steve: Aging. We're all aging.

Leo: You know what I have to learn now?

Steve: No.

Leo: How to program an HP-42 in RPN.

Steve: Oh. Isn't it beautiful?

Leo: I have no need for this at all. In fact, as soon as I got it realized there's a $10 software version for the iPhone. But this is really cool, from Swiss Micros.

Steve: Yeah.

Leo: And there was no tariff. I don't know how they did that. I got the DM-42n.

Steve: Yeah.

Leo: You have a 41, I think.

Steve: And I can see that it's got that cool pyramid on. And as you know, every time you turn it off you get a different graphic that it leaves...

Leo: This is a fractal, yeah. So somebody wrote a fractal in the thing.

Steve: And there's your register stack.

Leo: Oh, there's a QR code. I don't know what that's for. Don't scan that, kids. No, I think it's probably Swiss Micros. There's another program.

Steve: There's a grid of weird little 3D squares.

Leo: I think these are fractals, is my guess. Oh.

Steve: Oh.

Leo: I got an owl.

Steve: That's the wise old owl.

Leo: This is all I know how to do right now, Steve. But it is really beautifully made. You inspired me.

Steve: It is, yes.

Leo: And I thought, it's a fetish object. I have no use. I have computers. I have spreadsheets. I don't need this. Nor do you. Unless every once in a while you want to hard code...

Steve: I pick it up all the time.

Leo: Do you, for programming?

Steve: I absolutely do, yes.

Leo: Converting your hexadecimal to...

Steve: Yeah, I mean, and like I've got how many servers will fit on the edge of a pin, or the head of a pin, like, you know, things are important when you're, you know...

Leo: This is basically not very intuitive. I have got to read the manual. I never used one of these.

Steve: No, well, it's daunting; right? I mean, like because all the buttons have multiple functions, and there's like a programming mode. Oh, there's also like a configuration mode and all kinds of - there's a lot there.

Leo: I did set the time and date. I was able to do that.

Steve: Ah, that's nice. So it's no longer set to Swiss time.

Leo: Right. And it has a - this one is the newer one with a USB-C connector, so it can actually - I think the processor speeds up when you plug it in a little bit.

Steve: Oh, I've got USB-C.

Leo: You might have a micro. Do you have C? Okay.

Steve: Yeah, I do have C.

Leo: I was really pleased. It came within a few days. You talked about this last week, and it came in time for the next episode.

Steve: Well, and it slipped under DHL's new $800 minimum; right?

Leo: That's what happened. Because it was delivered by DHL. And it's the de minimis exception means it was not tariffed, it was under - it's only - Lisa said, "What did you buy from Switzerland for 300 bucks?" A calculator, honey.

Steve: Today. What did you buy from Switzerland today? Yesterday [crosstalk].

Leo: Yeah. Actually it was more like today. Yeah, it was. Anyway, what's coming up? Speaking of today on Security Now! today.

Steve: So while researching an interesting piece of security news, which we're actually going to get to next week, I strongly suspect, I stumbled upon a feature that we all have, we who have Windows 10 and 11, which is now the majority of the "we," that I thought, you know, we've never talked about this. I had forgotten that it was there. And it is - the more I looked at it, Leo, the more impressed I got. And our listeners know nothing Microsoft does recently really winds me up. I mean, you know...

Leo: No, it's true.

Steve: It doesn't. I am infatuated.

Leo: Oh.

Steve: I am so impressed...

Leo: My.

Steve: ...with the design of this. And everyone's going to know by the end of today's podcast about Windows Sandbox. Windows from 19 whatever, was it 1903, so very early on in Windows 10, Windows 10 acquired a stunning technology, which allows for another version, another instance of the Windows OS to be launched. And not like a VM, but like an app. That means they did everything about this right. It's in our Windows. I'm going to show everybody how to find it and turn it on because it's not on by default. That's why nobody knows about it. But it is a true security sandbox that allows you to run code that might be sketchy, download files you're not sure about. You could use Tor in it, and surf the 'Net. And when you close the sandbox, all trace of what you did is gone. And it launches in seconds as opposed to a VM that, you know, basically bringing up a whole new version of the OS.

Anyway, I'm excited to share with our listeners all of the features that it has, and also some of the why I am in love with this thing because, I mean, and I mention lower down in the podcast that this is maybe the first time that I've been envious that my other machine is still on Windows 7 and doesn't have this. I mean, I've been like, eh, 10, it's three digits higher than 7, you know. Otherwise, you know, who cares? But it's like, I want this. So anyway, but we're going to get there. We're going to first talk about enabling Firefox's tab grouping. The recalled Recall rerolls out. The crucial - you sent me a text, I think it was maybe Tuesday afternoon, about how the CVE program came very close to dying last week.

Leo: That would have been a shocker.

Steve: Oh.

Leo: I mean, that really would have been devastating.

Steve: Actually, it's not just that we wouldn't have had numbers for the podcast. It turns out it's actually crucial to, like, the whole management of vulnerabilities worldwide. And it almost went away. China has confessed, actually officially said, yeah, that was us, hacking the U.S., which they blame our stance on Taiwan.

Leo: So it's our fault.

Steve: So it's really your fault for making us do it to you.

Leo: We had to hack you, yeah.

Steve: That's right. CISA says what Oracle still refuses to. We've got brute force attacks on a very rapid rise. A very worrisome Python package which has a hard time not being a 9.8 CVS score, or CVSS...

Leo: CVE. CVS is the drugstore. CVE is the - yes.

Steve: Right. Oh, that's right, yes. Thank you. Also the CA/Browser Forum has passed the short-life certificates measure.

Leo: Oh, nuts.

Steve: We're going to revisit that. Maybe, well, certainly for the last time until it gets really bad. We have a few years left. But it's certainly not anything that anybody's going to be able to ignore any longer and hope doesn't happen. A wonderful crosswalk hack hit Silicon Valley last week. Android had the strangest announcement about that forced restart feature.

Leo: Yeah.

Steve: Anyway, we'll have some fun with that. Also we're going to look at how the EFF is never happy, but especially now about Florida. Some interesting research into ransomware payouts. And for Security Now! 1022, for not the last podcast, we squeezed a lot of podcasts into this month, Leo, because we started on the first. We started on April Fool's Day, which means we get one more podcast in April before we have to switch over to May. So podcast number 1022, Windows Sandbox.

Leo: Love it.

Steve: For the 22nd of April. And we're going to have a lot of fun in the next, what, about 12 or 13 hours that we'll be doing this.

Leo: Ah ha ha. Don't forget the Picture of the Week, also coming up.

Steve: Oh, it's actually a great Picture of the Week, yup.

Leo: It's a doozy. And I've got our caption of the week on my Swiss Micro calculator. "Don't panic," it says. "Don't panic."

Steve: We know where that came from.

Leo: Yeah. Stay tuned. We'll probably give you reasons to panic, actually, coming up. All right, Steve.

Steve: So you've not seen this picture. We have the caption "Why we will never have perfect security."

Leo: Okay. I'm going to scroll up right now, and I shall see the picture. Okay. You'd better describe this. I'm not sure I get it.

Steve: Okay. So...

Leo: Well, the door is open. Is that the joke?

Steve: And what puts the perfect punctuation on is the book that was used to hold the door open...

Leo: Oh, I missed that part. The CISSP Exam Book.

Steve: So for those who are not seeing the...

Leo: I missed that. That's hysterical.

Steve: So what we have is an endeavor to create a secure environment in the Security Operations Center of some facility somewhere. We have a door clearly labeled SOC, Security Operations Center. And underneath it, it says "Please Knock" because otherwise you ain't getting in.

Leo: Oh, "Access is Restricted." There's a whole sign that says that.

Steve: That is right. It's over to the left. It says "Security Operations Center. Access is Restricted." To enter, blah blah blah.

Leo: You need a card key or a...

Steve: Yes.

Leo: Yes.

Steve: There is an automated lock, and we have a card reader which is doing that. And, you know, you cannot get in. And pardon me, this thing is making noise.

Leo: You know what I think is the issue is I think there's no bathroom in the Security Operations Center. And so when you need to go, you need to prop the door open is my guess.

Steve: Yup. That's exactly right. So essentially what happened was - well, and I did also want to note that there is an electronic card key reader to the side. So, I mean, these guys are clearly serious about the security.

Leo: Oh, yeah. Oh, yeah.

Steve: What we find, however, is the door has been propped open, which of course completely defeats all of this - the sign, the knock knock, the warnings, the electronic card scanner. You don't need any of that because the door's not latched. And as you noted, the icing on the cake is that the CISSP security operations training book, which is clearly well used, it's got like little flags in...

Leo: Oh, yeah, bookmarks, Post-its, oh, yeah, oh, yeah.

Steve: Yeah, yeah, yeah. I mean, somebody took some time with this thing and decided, well, what the heck. What book is handy that I can use to prop this open?

Leo: That's really the - that's the funniest part of all. That's great.

Steve: It is great. And I think you're right. Benito and I were talking about this just before the podcast. And we agreed that it was probably the case that the guy forgot his badge at home.

Leo: Yeah.

Steve: Or, you know...

Leo: Just had to run out.

Steve: ...the dog ate it. There was nobody in there who would let him back in, so for him to do his knock, knock, knock routine, and he had to pee. So, like, oh, well. And again, why we will never have perfect security. And, you know, the larger point here is that this will always happen. Right? I mean, it is the human factor which is always going to be the problem. You know, phishing is the way people get in now, is they, you know, they send a piece of email that looks completely reasonable, and in fact it got Troy Hunt, as we talked about a couple weeks ago.

Leo: I know, I know.

Steve: Troy got phished.

Leo: Yeah.

Steve: Okay. So that was our Picture of the Week. After hearing last week's note about Firefox tab grouping, and how I've been unable to get a pair of tabs to merge on my Firefox, which was updated to 137, which was the one, the version that was supposed to have it, a number of our listeners said "Uh, Steve, it's probably there, but just disabled." And sure enough...

Leo: Oh, it's in the about:config.

Steve: Yes.

Leo: Ah.

Steve: It's not that I didn't have it, it wasn't enabled. So for anybody else who wants it, because I've got it now, and it's nice, about:config, then in the address bar, or in the search of about:config, search for tabs.groups. That will return three entries, and the amount of time you must hover and hold the tab before they merge, before Firefox says, oh, this guy wants to do a merge. Then also browser.tabs.groups.enabled. Mine was set to false; it's now true. And browser.tabs.groups.smart.enabled. I don't know what the difference is, but I want my tabs to be smart so I enabled that, too.

Leo: But it's probably disabled by default. I mean, I don't think...

Steve: Yes. They were both - they were all - so what I believe is that when they talked about it gradually rolling out, what's gradual, what's...

Leo: Is turning it on, yeah, yeah.

Steve: Exactly. So all the code is already there in everybody's Firefox 137 and later. And without you doing anything, they'll, I don't know what, give you a little popup and say, hey, you could try this now. I don't know how they're going to tell people. Or maybe it'll just start working, and people go, oh, look at that, I hovered my tab, and they joined instead of sliding past each other. Anyway, it's there for anybody who's interested.

And so I disabled my Tree Style tabs add-on, which is what was giving me vertical tabs in a nested hierarchy. You know, I shed a tear for the tree, the lack of a tree architecture. And I don't really use it that much, but sometimes I'll put stuff under a tab and then close that tab. But Firefox has the same thing. It allows you to assign a group name, and you can click that group name in order to collapse all the tabs under that. So I can easily see this solving people's problems. I also like a little more density, and I poked around a little bit, and you could get into a custom CSS style sheet, which is used to format the, you know, so-called "chrome" around the edges of the browser. But I thought, ah, I'm just going to see if I get used to, you know, change is hard; right? So I'll just get used to it and, you know, see if I get used to it and work with what's the default.

But anyway, so, and I didn't mention that I also am using Firefox's native vertical tabs. So I've got vertical tabs, and now I've got tab merging that lets me create groups that are named, and you can set the color and do different things. So, yeah. We've got it. And it's all there for anybody who wants it.

This news would have made it into last week's podcast, except that last week already broke the record, world's record, for the longest Security Now! podcast ever, which is why I was joking earlier about, you know, this one being maybe 12 or 13 hours. No. There was no room available to talk about this. And just so that everyone knows, because I did get some feedback from people saying, Gibson, three hours, really? Come on. I recognize that three hours is a lot of everyone's life. And I did hear your pushback. So that was just, you know, it wasn't an intentional marathon. We're not deliberately, you know, extending the length of the podcast. It was just that there was a lot to talk about. So anyway.

The original announcement, which, as I said, I would have gotten to last week, was of the release of a new Windows 11 to the Release Preview Channel, which was made on April 10th. And, now, that was for Build 26100.3902. But that release apparently had a few issues that cropped up pretty quickly because Microsoft has since updated it to .3909 from .3902, and that was on last Friday the 18th. And after all, quick updates are what you expect. That's the inherent nature of Preview Releases. You know, things are going to be discovered due to wider deployment, and then they're going to get fixed for everybody.

So anyway, because Microsoft now clearly recognizes that their Copilot+ "Recall" technology, which created quite a stir when they tried to do this a year ago, is a big deal, and does really represent a huge change to the operation of Windows. It was the first new feature that they noted in their Preview release notice. Once Recall makes its way into the production releases, I'm sure it'll come up again. There'll be some press coverage about it. We'll probably take another look at it, you know, as will the entire Windows 11-using world. But as Microsoft promised last year, when Recall is initially released, it will be disabled by default. So that's big change number one.

And it will, you know, I'm sure they're going to be telling everybody, oh, don't worry about it. It's secure. It's encrypted. You know, your privacy comes first, blah blah blah. But it will be an opt-in feature of Windows 11, at least for now. We know that when Microsoft wants, really, really, really wants you to have something, like Xbox for some reason - I don't own an Xbox, but I've got it in the menu - then you're going to have it, whether you want it or use it or not. Anyway, it's there. I wanted to let people know that it's on its way, and that what we learned about what they would be doing with it has turned out to be the case. They understood that it's not something that can be opt-in and that everybody gets without question.

And also, as I mentioned, Leo sent me a text Tuesday after the podcast last week. For a few days last week, it appeared that the incredibly important and extremely useful Common Vulnerabilities and Exposures program that's operated by the Mitre Corporation, and always has been, and has always also been funded by DHS, the U.S. Department of Homeland Security, might become unfunded. And people were talking about it getting shut down. The entire security industry breathed a collective sigh of relief with the news that CISA found some loose change somewhere, enough to keep it going for another 11 months.

Last Wednesday, under their headline "CISA extends funding to ensure 'no lapse in critical CVE services,'" BleepingComputer wrote the following. They said: "CISA says the U.S. government has extended Mitre's funding to ensure no continuity issues with the critical Common Vulnerabilities and Exposures program." CVE, not as Leo corrected me, CVS, which is the pharmacy. That's different. The U.S. cybersecurity agency told BleepingComputer: "The CVE Program is invaluable to the cyber community, and a priority of CISA. Last night, CISA executed the option period on the contract" - which I guess was always there, but still we were all brought to the brink - "to ensure there will be no lapse in critical CVE services. We appreciate our partners' and stakeholders' patience.

"BleepingComputer," they wrote, "has learned that the extension of the contract is for 11 months. The announcement follows a warning from Mitre Vice President Yosry Barsoum that government funding for the CVE and CWE programs was set to expire today, April 16" - when this all happened last week - "potentially leading to widespread disruption across the cybersecurity industry. Barsoum said: 'If a break in service were to occur, we anticipate multiple impacts to CVE, including deterioration of national vulnerability databases and advisories, tool vendors, incident response operations, and all manner of critical infrastructure.'" I mean, this was a big deal.

"Mitre maintains CVE, a widely adopted program that provides accuracy, clarity, and shared standards when discussing security vulnerabilities" - and, you know, it's a staple for this podcast; right? - "with funding from U.S. National Cyber Security Division of the U.S. Homeland Department of Security.

"After publishing our story," wrote BleepingComputer, "Mitre shared the following statement with BleepingComputer. 'Thanks to actions taken by the government, a break in service for the Common Vulnerabilities and Exposures (CVE) Program and the Common Weakness Enumeration (CWE) Program has been avoided. As of Wednesday morning, April 16th, 2025, CISA identified incremental funding to keep the programs operational. We appreciate the overwhelming support for these programs that have been expressed by the global cyber community, industry, and government over the last 24 hours. The government continues to make considerable efforts to support Mitre's role in the program, and Mitre remains committed to CVE and CWE as global resources."

Leo: Mitre does. I don't know what CISA's planning. But okay. They've cut them way back.

Steve: We don't know, I mean, I don't know, I don't think anybody knows, to your point, Leo, which is a good one, what CISA is today.

Leo: Right.

Steve: We know what CISA was last year.

Leo: Right.

Steve: And we've been singing CISA's praises for years and been very impressed with CISA. Now, you know, as is the case with a lot of what's going on in Washington, we just need to wait and see.

Leo: Well, not to mention the shameful fact that the President has asked the Justice Department to investigate Chris Krebs...

Steve: Yes.

Leo: ...former director of CISA, for the sin of saying that the election in 2020 was...

Steve: Was the most successful or the most secure election we've ever had in the U.S., which it was.

Leo: Now, in all likelihood that investigation's not going to lead to anything. Right? It's just BS. But it's still kind of terrifying that that can happen.

Steve: It's the intersection of politics with our technological world. And...

Leo: Right. And the problem is, security doesn't care about politics.

Steve: No.

Leo: You know, the bad guys are going to do what they're going to do. And if we don't fund the defense, we're going to have trouble.

Steve: Yeah. So BleepingComputer said: "Before CISA's announcement, a group of CVE Board members announced the launch of the CVE Foundation. So this is part two of this news, a non-profit organization established to secure the CVE program's independence in light of Mitre's warning that this U.S. government might not renew its contract for managing the program.

"Mitre said in a Wednesday press release: 'Since its inception, the CVE Program has operated as a U.S. government-funded initiative, with oversight and management provided under contract. While this structure has supported the program's growth, it has also raised longstanding concerns among the members of the CVE Board about the sustainability and neutrality of a globally relied-upon resource being tied to a single government sponsor.'"

Leo: That's a very good point. It shouldn't be.

Steve: Single, right. And they said: "Over the last year, the individuals involved in the launch have been developing a strategy to transition the program to this dedicated foundation, eliminating 'a single point of failure in the vulnerability management ecosystem' and ensuring 'the CVE Program remains a globally trusted, community-driven initiative.'"

Leo: So this was a wakeup call, is what it was.

Steve: Yeah, exactly. And, you know, it was a good thing, too, because we haven't lost continuity. We get 11 months, which should be plenty of time. So, and BleepingComputer finished, saying: "While the CVE Foundation plans to release further information about its transition planning in the coming days, the next steps remain unclear, especially considering CISA has confirmed that funding for Mitre's contract has been extended. The European Union Agency for Cybersecurity (ENISA) has also launched a European vulnerability database (EUVD), which 'embraces a multi-stakeholder approach by collecting publicly available vulnerability information from multiple sources.'"

Okay. So first of all, it is difficult to imagine a world without some common, uniform system for ranking the dangers and threats of vulnerabilities. Now, lord knows the U.S. government probably obtains at least as much value and benefit itself from having this program in place as any other entity. CISA will provide, as we noted, an additional 11 months of federal funding to Mitre, making this a very valuable wakeup call for the rest of the industry, and giving it time to arrive at a non-government-funded alternative. Which takes us to the announcement of the CVE Foundation.

Leo: Oh, good.

Steve: Speaking of a non-government funded alternative, also last Wednesday the industry was treated to a press release from the newly formed "CVE Foundation." The press release read: "FOR IMMEDIATE RELEASE / CVE Foundation Launched to Secure the Future of the CVE Program." From Bremerton, Washington they sent: "The CVE Foundation has been formally established to ensure the long-term viability, stability, and independence of the Common Vulnerabilities and Exposures (CVE) Program, a critical pillar for the global cybersecurity infrastructure for the past 25 years.

"Since its inception, the CVE Program has operated as a U.S. government-funded initiative, with oversight and management provided under contract. While this structure has supported the program's growth, it has also raised longstanding concerns among members of the CVE Board about the sustainability and neutrality of a globally relied-upon resource being tied to a single government sponsor. This concern has become urgent following an April 15th, 2025 letter from Mitre notifying the CVE Board that the U.S. government does not intend to renew its contract" - after 25 years - "for managing the program. While we had hoped this day would not come, we've been preparing for this possibility.

"In response, a coalition of longtime, active CVE Board members have spent the past year developing a strategy to transition CVE to a dedicated, non-profit foundation. The new CVE Foundation will focus solely on continuing the mission of delivering high-quality vulnerability identification and maintaining the integrity and availability of CVE data for defenders worldwide.

"Kent Landfield, an officer of the Foundation, said: 'CVE, as a cornerstone of the global cybersecurity ecosystem, is too important to be vulnerable itself.'" I love that. He said: "'Cybersecurity professionals around the globe rely on CVE identifiers and data as part of their daily work, from security tools and advisories to threat intelligence and response. Without CVE, defenders are at a massive disadvantage against global cyber threats.' So the formation of the CVE Foundation, they wrote, marks a major step toward eliminating a single point of failure in the vulnerability management ecosystem and ensuring the CVE Program remains a globally trusted, community-driven initiative. For the international cybersecurity community, this move represents an opportunity to establish governance that reflects the global nature of today's threat landscape.

"Over the coming days, the Foundation will release more information about its structure, transition planning, and opportunities for involvement from the broader community. For updates or inquiries, contact info@thecvefoundation.org." So that's the URL: thecvefoundation.org. So it exists. And depending upon how things look 11 months from now, and maybe even so, I mean, certainly given the current administration's feeling about, you know, waste, fraud, and abuse, if there is a foundation willing to take this over, I'm sure it's going to be cut loose. So...

Leo: This mirrors what's been going on with the Internet since its inception. You remember when it was one guy at UCSD, Jon Postel, who would assign you your IP addresses? And then it was IANA, and then IANA became a non-governmental organization. ICANN became non-governmental. The Commerce Department used to fund it, used to run it, and then released it to the world. Because we invented it here. So initially we did it.

Steve: It was originally under the auspices of DARPA.

Leo: Right.

Steve: The Defense Advanced Research Project Agency.

Leo: Yes. So, I mean, this is just - this is a natural evolution.

Steve: Yeah.

Leo: It's good we had this little wakeup call. It's good that they did not defund it because there would have been an interregnum in which we didn't have any CVEs assigned. That's more than just assigning a number. I mean, right? It's important.

Steve: Yes. Yeah. It would be difficult. It's that there is an agreement about where these numbers come from. There are - and if you ever look at the actual NIST database, a vulnerability is broken down into a whole bunch, essentially sort of a demographic of the vulnerability. You know, and there are official designators for each different category that the vulnerability falls into. I mean, it's odd because it's like oxygen. You know? We breathe it in.

Leo: We've always had it, yeah.

Steve: We take it for granted. And it's like, what would we have if there was no way of saying - well. And I was going to say, to finish that thought, no way of objectively evaluating how bad a problem was because many people, you know, jump on a 9.8.

Leo: Right.

Steve: It gets their attention. They know...

Leo: Right.

Steve: ...they have to fix this.

Leo: It's serious, right.

Steve: And if it's a 4.2, it's like, okay, we'll wait till next month, you know, because my shoe won't fit.

Leo: And you don't have to have a memory to understand what it would be like if there weren't a central naming authority because that's how virus names, and every security researcher has a different name for viruses. Same thing with threat groups; right? Everybody's Fancy Bear and...

Steve: And it's, yes, it's a mess.

Leo: It's a mess.

Steve: Yes.

Leo: You need a centralized somebody that says this is what we're going to call it. We all agree; right? This makes sense.

Steve: Yup. Yup.

Leo: Do you want to continue on?

Steve: It would also make sense...

Leo: Yes, I knew, I guessed that, yes. Let's take a little break. We will come back with more of Steve and this fabulous show. We're so glad you listen, and I'm so glad Steve continues every week to put - he puts so much work into this. And I'm very grateful, Steve, because it's not only our most-listened-to show, it is also the most important show we do.

Steve: Well, and I know my time is being well spent because I get so much feedback from our listeners. I sent out - we're now on the high side of 17,000 email subscribers. It was 17,097 last night received the advance notice and these notes and the Picture of the Week and so forth. And I get feedback from people saying, hey, you know, you've already got your tab merging in Firefox or whatever. So it's a resource for me.

Leo: Great resource.

Steve: And I know it matters to people.

Leo: Yeah, I've always thought of what TWiT does really ultimately is as a user group. You know, it used to be you would go to your user group every month.

Steve: Yup.

Leo: And you'd learn about, you know, your Atari 800 or whatever it was. But now that computing is ubiquitous, there isn't a place people can go to, you know, share this knowledge. And so this is - that's what you are, and that's what this is becoming. I think it's very, very important. It's really what the whole network is all about. So we appreciate your support for what we do here.

Steve: Wouldn't be here if it weren't for you, Leo. So there.

Leo: We're all glad to be here. That's all I can say. And the alternative is much worse. I bemoan my age to Lisa, as I'm sure you do to Lorrie, and she says, "Well, consider the alternative."

Steve: Yeah. Yeah, I have a variation of that. I tell people, "My plan is to live forever. And so far it's working."

Leo: It is. And we're so glad. Ah, I knew you would get around to this eventually. Let's talk about China hacking us. Oop, you're muted. Muted.

Steve: Sorry.

Leo: There we go.

Steve: My fault. Yes, so The Wall Street Journal carried the news under their headline "In Secret Meeting, China Acknowledged Its Role in U.S. Infrastructure Hacks," and they gave it the subheading: "A senior Chinese official linked intrusions to escalating U.S. support for Taiwan." Right, uh-huh.

Leo: And there's a lot of other reasons; right? Tariffs and so forth.

Steve: Well, but of course the China hacking has been going on for quite a while.

Leo: That's true. Pre-tariffs, yeah.

Steve: So it's like, come on, folks, really? The Journal story said: "Chinese officials acknowledged in a secret December meeting that Beijing was behind a widespread series of alarming cyberattacks on U.S. infrastructure, according to people familiar with the matter, underscoring how hostilities between the two superpowers are continuing to escalate. The Chinese delegation linked years of intrusions into computer networks at U.S. ports, water utilities, airports, and other targets, to increasing U.S. policy support for Taiwan, the people, who declined to be named, said."

So the attribution of these attacks to state-sponsored groups, specifically "Volt Typhoon," has been officially substantiated, and we have further evidence of what seems to me like

a bizarrely intertwined and complex relationship between our two countries. You know, we talked last week about the offhand comment that I heard from somebody who was being interviewed on one of the Sunday shows, saying that, well, at some point China might decide to weaponize all the information that they have been, you know, absconding with from the U.S. And it's like, oh, I hadn't thought about that. That would not be good, either. So it's like I just wish we could all get along, but doesn't look like that's going to happen any time soon.

As one security news reporter wrote: "CISA has published an alert on the Oracle Cloud data breach before Oracle did, mainly because the company is still busy wordsmithing its way around the issue." CISA's Alert published last Wednesday was titled: "CISA Releases Guidance on Credential Risks Associated with Potential Legacy Oracle Cloud Compromise." So, you know, because Oracle hasn't said anything official, CISA is having to tiptoe a little bit; right? I mean, they just can't come out here and blast away at Oracle. So they're being as careful as they could be.

They wrote, in this announcement last Wednesday: "CISA is aware of public reporting regarding potential unauthorized access to a legacy Oracle cloud environment." It doesn't get any more kid gloves than that. "While the scope and impact remains unconfirmed, the nature of the reported activity presents potential risk to organizations and individuals, particularly where credential material may be exposed, reused across separate, unaffiliated systems, or embedded, for example, hardcoded into scripts, applications, infrastructure templates, or automation tools. When credential material is embedded, it's difficult to discover and can enable long-term unauthorized access if exposed.

"The compromise of credential material, including usernames, emails, passwords, authentication tokens, and encryption keys, can pose significant risk to enterprise environments. Threat actors routinely harvest and weaponize such credentials to escalate privileges and move laterally within networks; to access cloud and identity management systems; to conduct phishing, credential-based, or business email compromise campaigns. They may resell or exchange access to stolen credentials on criminal marketplaces and enrich stolen data with prior breach information for resale and/or targeted intrusion.

"CISA recommends the following actions to reduce the risks associated with potential credential compromise. And this is generic at this point." They said: "Reset passwords for any known affected users across enterprise services, particularly where local credentials may not be federated through enterprise identity solutions," which would otherwise make them secure. "Review source code, infrastructure-as-code templates, automation scripts, and configuration files for hardcoded or embedded credentials, and replace them with secure authentication methods supported by centralized secrets management.

"Monitor authentication logs for anomalous activity, especially involving privileged, service, or federated identity accounts, and assess whether additional credentials such as API keys and shared accounts may be associated with any known impacted identities. Enforce phishing-resistant multifactor authentication for all user and administrator accounts wherever technically feasible. And finally, for additional information for or on cloud security best practices, please review the following Cybersecurity Information Sheets." And they give their title: "CISA and NSA Release Cybersecurity Information Sheets on Cloud Security Best Practices."

And then for users they only have three points: "Immediately update any potentially affected passwords that may have been reused across other platforms or services. Use strong, unique passwords for each account and enable phishing-resistant multifactor authentication on services and applications that support it. For more information on using strong passwords," blah blah blah. And finally: " Remain alert against phishing attempts,

you know, referencing login issues, password resets, or suspicious activity notifications. Be very skeptical." And then they reference their phishing guidance called "Stopping the Attack Cycle at Phase One."

So that advice could hardly have been more generic. That doesn't mean it's not obviously useful advice. But it does mean that in the absence of any confession from Oracle, you know, that's about as definitive as anyone is able to be. CISA felt that they had to say something because Oracle was really being irresponsible. I mean, this has been a sad lesson. You know, while I doubt that Oracle's irresponsible behavior will hurt them in the very short term, no one who's involved in the security industry is likely to forget this. It really should cause everyone to wonder, if they will act this way, what else is their internal corporate and security culture likely to do? And so the question is, how can you trust them? And unfortunately, these days more than ever, trusting the suppliers of critical infrastructure is all we really have. And Oracle hasn't indicated, hasn't demonstrated that they deserve that trust.

And speaking of MFA, Multifactor Authentication, I wanted to share a recent useful and important and even thought-provoking piece from the security firm Rapid7. Their piece was titled "Password Spray Attacks Taking Advantage of Lax MFA." Now, of course, multifactor authentication we've talked about a lot. I've recently encountered, and this is the reason I wanted to point this out when this popped up again, because I've been encountering reports for the last few months of significantly increased brute force guessing attacks, known as often credential stuffing attacks now. I recall us taking a close look at some problems that McAfee had a number of years ago. And what stood out was that bad guys were just pounding away at their login pages while McAfee was apparently blissfully unaware that anything was going on outside.

And of course just offering multifactor authentication is not a guarantee of safety itself. We recently looked at Microsoft's mis-designed MFA system which was allowing massive multifactor authentication brute forcing - enough to bypass that "million guesses required" barrier which, you know, is presented by any random six-digit passcode. But the more factors that can be added without unduly inconveniencing the user, the better. And as we've also seen, being smart about the deployment of MFA or even, you know, the use of a backup email loop for confirmation, where, for example, connecting to any previously seen IP or carrying a known browser cookie can be used to shift the security of a login in the direction of increasing the user trust.

So instead of, like, always asking for an additional authentication factor, if the user has provided a username and password, and is connecting from an IP that they've previously authenticated themselves from, then, you know, let's cut them a bit of slick, you know, not requiring them to jump through so many hoops. Or if they're using a browser that has a secure cookie token that was previously issued under multifactor authentication, then, okay, clearly it's the same person coming back, requires some authentication, but don't make it too intrusive. So being smart about multifactor authentication makes sense.

So here's part of what Rapid7 wrote. They said: "In the first quarter of 2025, Rapid7's Managed Threat Hunting team observed a significant volume of brute force password attempts leveraging FastHTTP, a high-performance HTTP server and client library written in Go, to automate unauthorized logins via HTTP requests. This rapid volume of credential spraying was primarily designed to discover and compromise accounts not properly secured by multifactor authentication. Out of just over a million unauthorized login attempts we observed," they wrote, "the distribution of originating traffic sources is similar to that previously seen just in January of 2025." So they're saying they took a much larger multi-month sampling, but the demographics of the sources of the attempts did not shift.

Some of the most prominent nations serving as points of origin for these attempts are Brazil, interestingly, at 70%, the huge majority of tax. Brazil at 70%. Then it drops immediately, Venezuela at 3%, Turkey at 3%, Russia at 2%, Argentina at 2%, and Mexico at 2%. So something's going on in Brazil that they've got 70% of all the attacks, and then the rest are much more widely distributed. May just be the bots that are, you know, the nature of the routers that are infected, and also good Brazilian bandwidth connections for those entities.

Anyway, they wrote: "Rapid7 has consistently highlighted multifactor authentication as a primary concern across several threat research reports. By the midpoint of 2023, data for the first half of the year showed that 39% of incidents our managed services teams responded to had arisen from lax or lacking multifactor authentication. Our 2024 Threat Landscape blog highlighted that remote access to systems without multifactor authentication was responsible for more than half, 56% of incidents as an initial access vector, the largest driver of incidents." So again, remote access to systems, no multifactor authentication, more than half the time, 56% of the time, that's how the bad guys are getting in.

"The third quarter of 2024," they wrote, "saw 67% of incident responses involving abuse of valid accounts and missing or lax enforcement of multifactor authentication." They wrote: "This total sits at 57% for the fourth quarter 2024, in part because of a 22% increase in social engineering." So that's on the rise, as we've been seeing and talking about. "Even without pausing to consider user agent-centric password spraying, this is a potentially dangerous combination for organizations not making the most of MFA-centric protection. If the brute forcing doesn't get you, a social engineering campaign might just do the trick," is what they said.

"Why MFA Matters, and the consequences of 'We'll Set It up Later.'" They wrote: "MFA is a key component of an overall Identity Access Management (IAM) strategy. If you're not making use of it, then your overall defense is weakened against many of the most common threats out there, including phishing. The very best password you can muster is made entirely redundant if your employee hands it over to a phisher, whether via a forged website or a social engineering attack. One way to mitigate against this is to use a password manager, which will only automatically enter your details on a valid website." We were just talking about that recently, Leo.

Leo: Yup.

Steve: And the benefit of requiring an exact domain name match, which, you know. And in fact it was Troy, right, who did not get the domain name match and said, well, you know, that happens.

Leo: I'll give them the MFA anyway, yeah.

Steve: Yeah, exactly.

Leo: That happens.

Steve: "But," they wrote, "what happens if your password manager's master password is compromised, and all the logins contained within are exposed? One of the best ways to

address this additional headache is MFA for all your accounts, including your password manager." And there I'll just say it's a reason to have...

Leo: Worked for Troy.

Steve: Yeah. Well, and it's a reason not to put MFA in your browser. Again, better to do it than not use it at all. But it is better to use it on a separate device. You know, that's where mine is. It's in my phone, which is always right next to me.

Okay, what about malware? They wrote: "Do you know what malware, password stealers, and keyloggers love more than anything else? Grabbing all those passwords stored in web browsers, or (in more serious cases) plain text files on the desktop" - do people still do that? Probably - "and email drafts. Do you know what they don't like? Having all of those perilous passwords protected with an additional layer of security. MFA could make the difference between compromise and data exfiltration versus a last-minute save and a security training refresher."

And finally, credential stuffing: "An unfortunate by-product of years of data breaches, often with phishing as the launch pad, roll-ups of new and ancient login details published online are a constant threat. It's worth noting that it isn't just your current employees who could be on these lists. Ex-employees with still-valid credentials are a cause for concern, too."

So they finish with: "Here are some steps you can take now to improve your security posture and mitigate risk from attacks like these, courtesy of Rapid7's experts: Number one, implement multifactor authentication across all account types, including default, local, domain, and cloud accounts, to prevent unauthorized access, even if credentials are compromised. Use conditional access policies to block logins from non-compliant devices or from outside defined organization IP ranges." Conditional access policies, meaning something else, some other block than just login credentials. And I'll have something to say about that in a second, something that I myself do. "Third, ensure that applications do not store sensitive data or credentials insecurely, for example, plaintext credentials in code, published credentials in repositories, or credentials in public cloud storage."

Next: "Audit domain and local accounts as well as their permission levels routinely to look for situations that could allow an adversary to gain access by obtaining credentials of a privileged account. These audits should include whether default accounts have been enabled, or if new local accounts are created that have not been authorized. Follow best practices for design and administration of an enterprise network to limit privileged account use across administrative tiers."

And this is really more important than it gets enough attention. The idea being, you know, an audit isn't something that you have to do, but it is clearly something that you should do. You know, you don't know what you don't know unless you do an audit of accounts. And there are so many instances where employees leave with their credentials, and we've covered these situations on the podcast, where they're disgruntled, they wait a week or two, then they log back in and do some damage, or get up to some mischief that they wouldn't be able to if their account had been deleted the moment they walked out the door, as should be the case, or maybe even beforehand.

Also: "Regularly audit user accounts for their activity, and deactivate or remove any that are no longer needed." It's a good point. You know, look at accounts that haven't been used in a long time. I'm sure that all of our more sophisticated users will often, for example, sort a directory by date and look at the really old stuff that hasn't been touched

in a long time and say, hey, you know, I don't need this any longer. Let's get rid of it. So lack of use is another really useful and easy-to-detect indicator.

They said also: " Wherever possible and aligned with business requirements, disable legacy authentication for non-service accounts and users relying on it. Legacy authentication, which does not support MFA, should be replaced with modern authentication protocols." And here, you know, Microsoft gets heat for having implemented insecure authentication originally, back when it really wasn't a big deal, back when it was only for local networks because no Internet existed back in the LAN Manager days. So security just wasn't an issue. Unfortunately, it has carried on into today's world with the Internet where security is an issue, and for the sake of backward compatibility they just, to their credit, they don't break old stuff. Unfortunately, they don't break old insecure stuff, either. So legacy can be a problem.

And finally they said: " Applications may send push notifications to verify a login as a form of multifactor authentication. Train users to only accept valid push notifications and to report suspicious ones."

And they conclude, saying: "You cannot go wrong with multifactor authentication. Imagine a scenario," they wrote, "where your network is under fire from a worryingly high number of brute force attempts from across the globe, targeting your insecure accounts until just one is compromised. Now imagine that same scenario where everything is blocked by default; regional restrictions are applied; logins from user agents are not allowed; and all your VPNs, your RDP, VDIs, and SaaS tools are secured with MFA.

"This may feel like an overreaction to what you may view as an attack that looks like an edge case. However, consider that ransomware groups, alongside more commonly found malware actors and publishers, will also find you a significantly harder target to break as a result of these countermeasures being put in place. Please don't end up in the unenviable percentage of organizations compromised due to missing multifactor authentication in our next threat research report." In other words, don't have your name among Rapid7's compromised companies.

They said: "There's no better time than now to think about building out a stronger security posture." And again, it's that "We're going to get around to it later" attitude. Just, you know, get it done. But all this amounts to is adopting a multilayered security approach. Never assume that any single protection will be sufficient. And a username and password is a single layer of protection. If it's possible to practically do more, do more.

Leo: Does Passkeys count as doing more?

Steve: Yeah. Passkeys is an absolutely, you know, another...

Leo: It feels unmultilayered; right? I mean, it's just one thing.

Steve: It is, but it is dynamic inasmuch as it is not subject to credential theft. So nobody can steal anything from the server because, very much like SQRL, Passkeys give servers no secrets to keep.

Leo: Yeah.

Steve: So if they have no secrets, they're not in danger of losing them. Some of the strongest security protections can be somewhat brittle and troublesome. I know that, Leo, you and I cannot login remotely to our SSH servers without a client having the proper private key to verify its identity.

Leo: Right.

Steve: Now, could that cause some inconvenience? Sure it could. But no way am I willing to expose an unmonitored SSH server that's only protected by a username and password, no matter how secure they might be. That's just not safe.

Leo: Come to think of it, it is kind of like Passkeys to use a certificate, you know, public key instead of passwords.

Steve: Yes.

Leo: It seems more convenient and easier, but it's more secure. That's good.

Steve: Yeah. It is, it is, yes, it's very secure.

Leo: Yeah.

Steve: And as another example, filtering some classes of remote connections by IP will mean that those filters, if you put filters in to only accept some types of remote connections by source IP, that will mean that those filters will break when IP addresses change. I had that happen to me, it was two weeks ago, when a cable modem died, and I needed to switch to another. My cable provider, Cox, was wonderful throughout the process, but I wound up with a never-before-seen residential IP address that was different from the one my previous cable modem had, and a great deal of my network infrastructure fell apart.

Leo: Oh, shoot.

Steve: But, you know, I was prepared for that. I had previously made notes of all the many places I had and I was using IP-based blocking or permission filters...

Leo: Smart, smart.

Steve: ...that would need updating, and I had previously arranged to be able to do that remotely in the event of a residential IP address change. Now, of course, IP-based permissions is only one layer of my security. But I just, I've said it before, I want to make sure everybody understands how awesomely powerful that layer is, so much so that it is well worth the hassle and a bit of brittleness where, you know, every three or five years or so my cable modem's IP may need to change. It doesn't happen often. But it does happen.

Leo: Right.

Steve: So anyway, I think that the ultimate takeaway from Rapid7's posting is to appreciate that there really are extremely determined, anonymous, and numerous attackers who are more or less continually pounding away, largely unmonitored, outside our gates. You know, we talk about monitoring our network. We don't really spend much time talking about monitoring the other side of our boundary, the other side of that barrier that's keeping the bad guys out.

And it is horrifying, if you look at, like, what's going on out there. You know, they couldn't - and it's not about you. That's the other important thing. They could not care less who you are. It's no longer reasonable to say, "Well, I'm nobody that anyone would want to hack." They don't know that until after they're in. And then, once they're in, the least they will do is arrange to establish persistence so they can mine crypto or use your bandwidth to increase their next DDoS attack. So, you know, I just - I can't stress it strongly enough. You don't want to be that kind of victim. But you do want to be a customer...

Leo: Oh, yes.

Steve: ...of TWiT's next sponsor.

Leo: Of BigID. You are - so prescient of you to realize that. Wait a minute, you just went full screen. What happened?

Steve: You're right, I did.

Leo: It's magic.

Steve: I gave it a thumbs-up, and it went full screen.

Leo: Well, I'm not touching anything. Both hands, see? I'm not touching anything. We'll see what happens. Steve? On we go.

Steve: So there's a Python library known as BentoML (B-E-N-T-O-M-L). Pretty popular. And as with pretty much anything "ML," the "ML" stands for Machine Learning. BentoML is a project over at PyPI which bills itself as "the easiest way to serve AI apps and models." Unfortunately...

Leo: Oh. Had to be an unfortunately.

Steve: That's why we're talking about it here.

Leo: Oh, well.

Steve: Since it also carries a CVSSv3 vulnerability and exploitability score of the difficult-to-attain 9.8, if you're unfortunate enough to be using v1.3.8 through 1.4.2, it may also be the easiest way to have your AI-related service taken over by bad guys thanks to the presence of a critical remote code execution vulnerability.

BentoML's documentation page explains that it's "A Unified Inference Platform for deploying and scaling AI models with production-grade reliability, all without the complexity of managing infrastructure. It enables your developers to build AI systems 10x faster with custom models, scale efficiently in your cloud, and maintain complete control over security and compliance." Sounds great. Except that apparently it's the bad guys who get to have the complete control over security, or lack of any. Since it seems pretty clear that we're on the brink of a new renaissance in AI-based security threats and vulnerabilities, I figured it would be worth taking a brief closer look at this one.

Here's what the security research group Checkmarx wrote. Checkmarx took a close look at BentoML. They said: "A critical Remote Code Execution (RCE) vulnerability, with a CVE (thank god for CVEs) 2025-27520 with a score, a base score of 9.8" - which is, you know, difficult to get - "they said has been recently discovered in BentoML, an AI service helper Python library found in PyPI. This flaw allows unauthenticated attackers to execute arbitrary code by sending malicious data payloads as requests and potentially take control of the server. While the advisory specifies versions from 1.3.4 through 1.4.2 as being affected, Checkmarx Zero's analysis indicates that this issue affects versions 1.3.8 through 1.4.2." In other words, fewer. "It is recommended that affected adopters upgrade to version 1.4.3 or later to repair the issue." And I will come back to why, or later maybe a bit of a question.

They wrote: "You are potentially affected by this issue if you use BentoML (either directly or indirectly) to receive and process machine-learning 'payloads,'" they said, "(which are serialized data structures), from untrusted sources. Since this is a primary purpose of BentoML" - in other words, that's what you use it for - "the presence of a vulnerable version of this library should be considered a significant indicator of actual risk." In other words, arranging to provide BentoML with a malicious serialized payload will not be difficult since that's what BentoML is designed to take in.

Okay. So Checkmarx wrote: "CVE-2025-27520 is a Remote Code Execution vulnerability found in BentoML, a Python library designed for creating online serving systems that are optimized for AI applications and model inference. The full GHSA advisory describes the vulnerability and exploitation, which we summarize here. The flaw, that originates from an insecure deserialization, enables adversaries to execute arbitrary code on the server by sending a specially crafted HTTP request. This issue exists because the deserialize_value function in the serde.py file deserializes input data without proper validation, meaning attackers can inject malicious payloads that trigger execution of arbitrary code when they are deserialized."

Okay. By now, any of this podcast's long-term listeners will perk right up when they encounter the term "deserialization" since we've previously encountered so many instances of deserialization gone bad. As we know, "serialization" is the process of taking a complex data structure and converting it into a stream of bytes, thus "serializing" it. So "deserialization" is the reverse process that takes as its input a previously "serialized" byte stream and hopefully returns the original complex data structure.

The reason we keep encountering security-related problems with deserialization is that the act of deserializing requires - here's another one of our problem words - requires the "interpretation" of the meaning of that serialized byte stream, and "interpreters" are notoriously problematic to get perfect, and any imperfection can too often be leveraged to create an exploitable vulnerability. What's even more unfortunate is that this is not -

and here it comes - not the first time the BentoML has had this 9.8 severity trouble. The NIST already had a listing last year in 2024 for CVE-2024-2912, to which it assigned the rarest of rare - oh, I was wrong, it's not a 9.8, it was a 10.0.

Leo: Ooh.

Steve: Uh-huh. And that's not surprising when a vulnerability disclosure describes the problem by writing: "The BentoML framework is vulnerable to" - get this - "an insecure deserialization issue that can be exploited by sending a single POST request to any valid endpoint." Meaning to the server. "The impact of this is remote code execution," they wrote.

So then Checkmarx writes of the newly discovered flaw: "This flaw is essentially a reintroduction of CVE-2024-2912, which had been previously fixed in version 1.2.5. Both CVEs deal with the exact same issue, an Insecure Deserialization vulnerability that can be exploited by sending an HTTP request to any valid endpoint to trigger remote code execution." At this point - this is me speaking - anyone using BentoML might reasonably question the wisdom of continuing to rely upon the developers of this package to keep them safe.

The Checkmarx guys wrote: "To exploit this vulnerability, the first step is to craft a malicious 'pickle.'"

Leo: Yes, well, that's the thing.

Steve: Beware the malicious pickle.

Leo: The malicious pickle.

Steve: That malicious pickle. They said: "A binary data serialization system commonly used with Python." Because it's pickled. They said: "This 'pickled' data payload contains Python objects that can contain executable code that gets run when the payload is deserialized for use by the application. Vulnerable versions of BentoML do not deserialize such payloads in a safe manner, meaning an adversary can send Python code which performs malicious actions including executing system commands under the authority of the Python application running on the server.

"In this case, an attacker can create a custom Python pickle object, for example, the Evil class, and override Python's magic method _reduce_ with a tuple that tells Python to run the os.system function. The _reduce_ method is used to specify how the object should be deserialized or serialized, and allows users to override default behavior with other meaningful actions." So this is, like, the power of Python being used against the server by the bad guys as part of the process of taking it over. They said: "By calling os.system, an attacker can trigger system commands during the deserializing operation, such as initiating a reverse shell connection to this machine, as shown in the provided proof of concept." So they provided the code to do this. And all the versions out there which are vulnerable are now known to be vulnerable.

Hoping to understand the sequence of events that caused a previously resolved and quite serious 10.0 problem to return, now as a 9.8, the researchers reconstructed the timeline of events. They wrote: "The vulnerability exists in BentoML versions 1.3.8 through 1.4.2.

If you are running a version within this range, you are affected. The advisory reports versions as early as 1.3.4 are vulnerable, but Checkmarx Zero analysts determined that the vulnerability actually re-emerged in version 1.3.8. Looking at commit 045001C3, we found that a previous security fix, originally introduced to address CVE-2024-2912, had been removed." Which, you know, why? Why was a previous security fix removed?

They said: "This missing code was specifically implemented to prevent this exact deserialization vulnerability now tracked as

CVE-2025-27520." They wrote: "So, first, the original vulnerability finding was reported as CVE-2024-2912. It was patched in version 1.2.5. The fix was later removed in version 1.3.8."

Leo: What?

Steve: "The same issue resurfaced and was reported again, now as CVE-2025-27520, and it has now been re-patched in version 1.4.3." So as I noted before, without some very clear accounting - and accountability - for these events, given the potential consequences of this library's direct exposure to the Internet so that a single HTTP POST query is all that's needed to completely take over remotely a system, anyone using or considering the use of this library would be well advised to proceed with extreme caution.

On the off-chance that any of our listeners might be affected by this, I've included the link to Checkmarx's posting and analysis. There isn't anything, you know, really tied to machine learning or AI about this. It just appears to be a very problematic Python library that appears to need better development management. We all know that mistakes can happen. That's the nature of the game. But if they're forgiven, they should be followed by some learned lessons. So let's hope that has happened here this time. You know, maybe people came and went. Different people are now running, I mean, without getting way closer, without being on the inside, it's impossible to understand how this happened.

But I would argue, before using this, someone should obtain an understanding of what happened and some reason to feel assured that it won't happen again. Or maybe use a version you know is not vulnerable, and very carefully scrutinize moving forward because it was by moving forward in the past that this problem was reintroduced. You don't want to have that happen to your server and get compromised as a result.

Okay, Leo. Now, this is where I need to just take a deep breath.

Leo: Okay. I'm going to have something that's going to wind you up a little bit right before you do this story.

Steve: Uh-oh.

Leo: This just released, this just out from Google, their privacy sandbox. They have decided to permanently change their policy. "We have made the decision to maintain our current approach to offering users third-party cookie choice in Chrome and will not be rolling out a new standalone prompt for third-party cookies." Remember they were going to phase out third-party cookies. They say: "After consulting" - remember they put it on hold a few months ago. "Now, after consulting publishers, developers, regulators, and the ads industry, we decided, yeah, I guess you need third-party cookies."

Steve: Consulting, in other words, consulting the people who pay our bills.

Leo: Yeah, exactly.

Steve: The people who allow us to live comfortably in Silicon Valley.

Leo: Yeah. I guess this was to be expected.

Steve: Well, and of course, you know, the larger part of this, you're right, I mean, the privacy sandbox, we were hoping that the system that Google came up with, which was really good...

Leo: Which you liked, yeah.

Steve: Well, remember, what it did was it transferred the responsibility to the user's browser.

Leo: Right.

Steve: The browser became the thing that was selecting ads. So it knew about its own users' historical browsing and was able to select meaningful ads on behalf of the user. I mean, it was a beautiful solution. But no, we can't have nice things, Leo.

Leo: But no.

Steve: No.

Leo: All right. Now that I've wound you up a little bit, go ahead.

Steve: Well, yeah, thank you.

Leo: Now you can get wound up a little bit more.

Steve: I guess I'm prepared. I'm prepared for another disappointment. On April 4th at 12:30, Ballot SC-081v3 was posted, and voting began. Sixteen minutes later, at 12:46, Chris Clements posted: "Google votes Yes on Ballot SC-081v3." The next day, Nick France posted: "Sectigo votes Yes on Ballot SC-081v3." That was followed two hours later by Apple's Clint Wilson posting: "Apple votes Yes on Ballot SC-081v3." The next day, Corey Bonnell posted: "DigiCert votes Yes on Ballot SC-81v3." And the day after that, Ben Wilson chimed in with: "Mozilla votes Yes on Ballot SC-081v3."

When the voting had ended, of the 30-member Certificate Issuers, 25 had voted yes, and no one, not one, voted no, though there were five abstentions. Of the four-member Certificate Consumers (Apple, Google, Microsoft, and Mozilla) all four voted yes. So what was it that just happened, essentially unanimously passing? This ballot was the formal adoption of a slightly toned down version of the quite aggressive certificate lifetime shortening proposal first made by Apple's Clint Wilson in October last year that set my hair on fire.

We talked about it at the time as I shook my head in bemusement. I don't understand it, and I probably never will, because the proposal appears to ignore all of the trouble that this will cause, while also conveniently ignoring the fact that 100% privacy-enforcing browser-side certificate revocation has finally been made to work. We have it. It's been working in Firefox. It's available to Chrome. It's in the public domain. Yet Clint's proposal just passed, and it did so handily. Clint's position is that nothing can be as certain as never issuing any certificates having long lives. That is, lifetime being short is impossible to have fail. It's impossible to get around. You don't need to rely on anything else. He's right. It's true. You can't disagree with him because factually he's correct. But I've seen no evidence to suggest that such an absolute level of certainty is warranted enough to offset the world of problems that this will cause.

Okay. So what just passed? Current certificate lifetime is a tolerable 398 days. So, you know, a month or two more than a year. We used to have 10 years. Then we had five years. Then we had two years. Now we have one year. So, you know, what we have now is effectively annual renewal and replacement with a little bit of slack. This 398-day maximum lifetime will be operable for any certificates issued before next March 15th of 2026. I will be reissuing all of mine the day before that because, on March 15th of next year, maximum lifetime will be summarily cut in half to 200 days, for no apparent reason that I'm able to divine. The year after that, on March 15th of 2027, lifetimes will again be cut in half to just 100 days. So this is essentially quarterly, requiring reissuance and renewal four times per year. Right? Because 90 days is a quarter plus 10 for some slack. At that point we either automate, or we spend our lives fussing with certificates.

And finally, in an apparent concession to some reality, the annual march to certificate lifetime extinction receives a two-year break since the final drop to just 47 days is deferred until March 15th of 2029. So we get two years for that final halving. But from then on, from March 15th of 2029, no certificate will be issued having a lifetime longer than 47 days. Why? I have no friggin' idea.

Leo: Well, that's the question. But what's the need?

Steve: There is no need. But that's the way it's going to be. And everyone has just signed onto that. What's clear is that anyone who is building any sort of device that needs to use public-facing certificates trusted by Chrome, Chromium, Firefox, and Safari is going to need to add ACME automation to their appliance, and they should start thinking about it sooner rather than later. For those running web servers, this shouldn't be any huge problem. There's a "win-acme" client for Windows servers, and actually there's about 10 different Windows solutions that I'm sure I'll be able to use. I haven't yet because I haven't minded, you know, going to DigiCert once a year and saying, hey, let's get me a nice updated certificate. I've got *.grc.com.

Well, you can't issue those from web browsers. You can only issue fully resolved certificates through ACME from web browsers. But you can use DNS. So I will have an automated solution that edits GRC's DNS for GRC.com so that I can receive a short, you know, 47-day eventually, certificate lifetime, you know, *.grc.com for GRC's servers and

other various domains. So I'm, you know, I'll solve the problem. Right? I don't have a choice. I'm mostly annoyed because no one has made, to your point, Leo, any clear case for why everyone needs to be so inconvenienced by this.

Leo: Is it because certificate revocation is broken?

Steve: No, it's fixed. We talked about it. With Bloom filters. It's working perfect.

Leo: Oh, that's right.

Steve: It's in Firefox. It's working. They've solved the problem. And remember the issue with OCSP, the Online Certificate Status Protocol, was that because your browser would reach out to the CA if a fresh OCSP certificate was not stapled to the TLS certificate that it had received, because of that there was some privacy concern. Your browser, your IP was reaching out to the Certificate Authority, saying is this certificate I've just received still good? So, okay. We will solve that with the Bloom filters, which we have. We now have CRL sets, you know, working certificate revocation on the browser side that is efficient, effective. It allows multiple times per day revocation. It's better than having the certificate last 47 days. It'll last a few hours before an updated set gets pushed out to all the browsers that are using browser side. So it's better than 47-day certificates.

And again, what happens if a DDoS attack lasts long enough so that when people come back and try to renew their certificates, they're unable to do so? Their certificates will expire, and their websites will go offline. This is brain dead. But it's what we're going to have. And I'll just say that the flipside is that ACME is already being widely used to dynamically generate the TLS certificates for 70% of all web servers worldwide. So, you know, it works. It's really only the laggard 30%, of which I'm a part, that needs to get with the program. And right now Let's Encrypt certs are renewed within a 90-day window. So, you know, not that much changes. Let's Encrypt will be bringing it down to shorter. Remember there was some concern about or some talk of a 10-day window? And so Let's Encrypt, the Let's Encrypt guys were saying, okay, we're going to gear up to, like, our whole infrastructure, so we're able to issue certificates 10 times more often than we do now. It's like, why? You need to, like, have a bigger budget? I don't get it.

But anyway, for what it's worth, it has happened. The industry said, okay, fine. Remember that Apple actually is the one who forced this by saying Safari would refuse to honor any certificates that had a longer life, that is, more time from not valid before to not valid after. Those are the two timestamps in certificates. So it's always possible to see how long a certificate could have lived. And Apple just said we're going to a year. If you don't bring your certs down to a window that we agree with, they're not going to be valid for any Apple properties. And so the industry said, uh, okay. We don't want to lose Apple.

Leo: I wish we knew why Apple thought this was important.

Steve: Well, the guy at Apple who is in charge of this happens to be a friend.

Leo: Oh. Well?

Steve: I'm thinking we should have Clint on the podcast.

Leo: Uh, yeah. You know him.

Steve: To talk with us about this.

Leo: Ask him, "Clint, why?" I mean, these are intelligent, rational people. He must have a good reason.

Steve: Yeah. I don't know what it is, but I think it would make a - he would make a great guest.

Leo: Okay.

Steve: So I will make that happen.

Leo: Send us the contact information, we'll reach out.

Steve: I'll make that happen.

Leo: Wow.

Steve: Crazy. Okay. I'm going to recover from that and have some coffee.

Leo: Okay, good idea. You're watching Security Now! with the unhappy Steve Gibson. He's not pickled, though. I just want to tell you that. He's not pickled.

Steve: I'm bemused.

Leo: He's bemused. He's bewitched, bothered, and bewildered. Steven Gibson, who has taken a deep breath and is ready to move on.

Steve: Ugh, yeah.

Leo: Well, you know, I mean, honestly, it's good for Let's Encrypt; right? I guess not everybody can use Let's Encrypt, though.

Steve: Yeah. And so, you know, any web browser, as I said, 70%, as we know, 70% of the Internet is now being secured with ACME automated Let's Encrypt certificates.

Leo: Right. Right.

Steve: So that tells you that, you know, people like...

Leo: Those people don't have to worry about it, yeah.

Steve: No. They don't. Except that they've got 90-day certs, and they're going to be cutting that in half. But okay, so what? You know, big deal.

Leo: Let's Encrypt will probably respond to that, yeah.

Steve: So, yes. It'll double the amount of traffic because Let's Encrypt will have to be renewing twice as often. But okay, fine. The 30%, there's a certain - certainly there's a chunk that are like me. I can solve the problem, I just haven't needed to. I've had other things to do. And so, okay, fine. I'm certainly not going to be issuing certificates four times a year. I will get my certificate, as I said, on March 14th of next year, so I get a whole year because you can't - because on March 15th maximum certificate lifetime drops to 200 days. Okay, big deal. But, you know, I'd rather have 400 days. Actually 398 days.

So again, that puts it off for another year before I have to worry about this. And during that time the ACME tools will mature more. There'll be better ways to solve the problem for Windows-based systems. The concern is things that don't easily automate, like as I mentioned, appliances that do need to be trusted, that have certificates in them, that need to be trusted by browsers that may not have had to automate for ACME yet. Those are going to - you're going to have to solve that problem.

Or the other thing that could occur is sort of interesting, too, because if there is a portion of the Internet which is using public certs because they've been easily available, but don't really need public certs, for example say that you had a telephone system, and the handsets all were using the PKI, the Public Key Infrastructure, and they had standard trust roots in them, and the equipment that the phone system was talking to was using, you know, certificate authority-issued certs, if this becomes too onerous, it would certainly be possible for the phone equipment supplier to become their own certificate authority. All they have to do is put their own trust root in the handsets, and then they issue certificates, and there's no one to tell them how long they could be. They could issue a 10-year certificate and just - and as soon as they supply the certs, their problem is over.

So we may see some fracturing of the public key infrastructure because it's been made too hard to use, because of what amounts to a special interest group of, you know, web browsers and servers that, for unfathomable reasons to me, want to have super short life certificates. I just, again, no one has shown me that we have a problem that we're trying to solve. But it's going to happen. They all just voted for it.

Okay. There is some fun news here, Leo. And I was thinking, I don't - I didn't put the link in here, so you can't bring it up, and I'm not sure, you don't want to get, you know, YouTubed or blacklisted or whatever it is happens.

Leo: YouTubed, I like that.

Steve: Last week TechCrunch carried the news of a pretty wonderful hack that hit the crosswalks...

Leo: Oh, we played this. It's okay, I'll play it. Yeah, yeah.

Steve: Okay.

Leo: I know what you're talking about, yeah.

Steve: ...that hit the crosswalks across the Northern California peninsula, commonly referred to as Silicon Valley. TechCrunch's headline was "Silicon Valley crosswalk buttons hacked to imitate Musk and Zuckerberg's voices." They wrote: "Audio-enabled traffic control crosswalk buttons across Silicon Valley were hacked over the weekend to include audio snippets imitating the voices of Mark Zuckerberg and Elon Musk. Videos taken by locals in Menlo Park, Palo Alto, and Redwood City in California show the crosswalk buttons playing AI-generated speech designed to sound like the two billionaires."

[CLIP] Hi, I'm Jeff Bezos. This crosswalk is sponsored by Amazon Prime with an important message, you know, please, please don't tax the rich. Otherwise all the other billionaires will move to Florida, too.

Leo: That's a new one. I hadn't heard the Jeff Bezos one.

Steve: Ah.

Leo: Here's another one. Here's another one. There's quite a few.

[CLIP] Amazon Prime.

Leo: Oh, no, we already had that one. Let's see. Here's another one.

[CLIP] Hi. This is Elon Musk. Welcome to Palo Alto, the home of Tesla Engineering. You know they say money can't buy happiness. And, yeah, okay. I guess that's true. God knows I've tried. But it can buy a Cybertruck, and that's pretty sick; right?

Leo: Had to bloop that one.

Steve: Uh-huh.

Leo: Let me see if I can find the Mark Zuckerberg one. There's quite a few. This is on X. I don't know. Do you have the story about how they did it?

Steve: No.

Leo: It turns out the PIN codes to protect these were 1234.

Steve: Uh, yup, yup. So what I was assuming. So TechCrunch wrote: "It's not clear why the sidewalk buttons were hacked, or by whom."

Leo: Because it's a good joke, that's why.

Steve: Oh. Yes. They said: "Palo Alto Online was one of the first outlets to report the hack, citing a Redwood City official saying that the city was 'actively working to investigate and resolve the issue as quickly as possible.'"

So TechCrunch finished their reporting by saying: "Audio-enabled crosswalk buttons are widely used across the United States to allow those with visual impairments or accessibility needs to hear custom audio messages that play for pedestrians to know when it's safe to cross a street. In a video from last year, physical penetration specialist and security researcher Deviant Ollam explains how audio-enabled crosswalk buttons can be manipulated, often by way of default-set passwords that have not been changed. Polara, the company that makes the audio-enabled crosswalk buttons, did not respond to a request for comment when contacted by TechCrunch on Monday."

Leo: This is the "Hacker Fun with Traffic Controls and Crosswalk Buttons" from Deviant Ollam. He explains the whole thing; right?

Steve: Yup.

Leo: These are for blind people, disabled people who can't see the lights, can't see the walk signs. They're audio walk signs, basically.

Steve: Right, right.

Leo: And they're everywhere in California. I don't know about elsewhere in the country, but they're everywhere in California.

Steve: So anyway, I wrote: "That's what we need more of, a bit of non-malicious, good old-fashioned techno-pranking."

Leo: Yeah.

Steve: You know? At the same time, that capability could have just as easily been used to produce extremely offensive audio messages instead of having some fun spoofing Zuck and Musk.

Leo: Pretty good impressions, too, I might add.

Steve: Yeah. And Bezos. Absolutely recognizable voices.

Leo: Yeah.

Steve: So anyway, a tip of the hat to the people behind that one, which, you know, we also have the benefit, as a consequence of this, of firmly closing whatever back door had been inadvertently left open, you know, thanks to this benign prank.

Leo: It's speculation. This is from Xeno Kovah on X, that the default password was 1234 on these. And I bet you, you know, some guy on a construction crew is installing it. He doesn't know to change the password.

Steve: Yup. Or somebody'll do it later; right? We need to leave the password so they can log into it and set it up. And whoever did that said, oh, I'll get around to it later. And they get around to it later. Not a good idea.

Leo: Get around to it. Everybody needs a roundtoit.

Steve: So last week, I noted that features similar to Apple's Lockdown Mode were expected to be announced during next month's Google I/O 2025. It appears that one of those forthcoming features could not wait. The features for Google Play services v25.14, dated last Monday, which was 4-14, April 14th, listed under "Privacy & Security" the following. It wrote: "Enables a future optional security feature, which will automatically restart your device if locked for three consecutive days."

Now, I'm not 100% clear, Leo, about what it means to "enable a future optional security feature." You know, the optional part I get. That's fine. I have no problem with that. But what exactly does it mean to "enable a future feature?" It's apparently now been enabled, which is why it's been listed. But if so, then how is that a future feature if it's already happened? It sounds like some change was made that we cannot actually use today, but we will be able to, optionally, in the future. In that case, who the hell cares? Why tell us anything about a future security feature that hasn't actually been enabled yet, even though it says it has been, you know, because we're still back here in the past? I don't know. I'm confused. But whatever it is, it's there, even if it's not really there. It's enabled, though you can't use it until the future.

Leo: Someday. Someday we'll all be able to use it.

Steve: Optionally; right. Okay. I wanted to share a write-up by the EFF over their "extreme unhappiness" over new legislation that's being proposed in Florida. And we'll understand why they're extremely unhappy. It's been my observation of the EFF that they are never happy. I mean, they're not happy about anything. You know, like, I mean, they're just so far out there. But, okay. We need them. I'm glad that we have a well-funded Electronic Frontier Foundation staffed by lawyers who know constitutional law. In this case, I don't know what Florida's thinking.

Here's what Florida said, and this is what the EFF wrote. They said: "At least Florida's SB 868/HB 743, 'Social Media Use By Minors' bill, isn't beating around the bush when it states that it would require 'social media platforms to provide a mechanism to decrypt

end-to-end encryption when law enforcement obtains a subpoena.'" They said: "Usually these sorts of sweeping mandates are hidden behind smoke and mirrors, but this time it's out in the open: Florida wants a backdoor into any end-to-end encrypted social media platforms that allow accounts for minors. This would likely lead to companies not offering end-to-end encryption to minors at all, making them less safe online. Encryption is the best tool we have to protect our communication online. It's just as important for young people as it is for everyone else, and the idea that Florida can 'protect' minors by making them less safe is dangerous and dumb.

"The bill is not only privacy-invasive, it's also asking for the impossible. As breaches like Salt Typhoon demonstrate, you cannot provide a backdoor for just the 'good guys,' and you certainly cannot do so for just a subset of users under a specific age. After all, minors are likely speaking to their parents and other family members and friends, and they deserve the same sorts of privacy for those conversations as anyone else. Whether social media companies provide 'a mechanism to decrypt end-to-end encryption' or choose not to provide end-to-end encryption to minors at all, there's no way that doesn't harm the privacy of everyone.

"If this all sounds familiar, that's because we saw a similar attempt from an Attorney General in Nevada last year. Then, like now, the reasoning is that law enforcement needs access to these messages during criminal investigations. But this doesn't hold true in practice. In our amicus brief in Nevada, we point out that there are solid arguments that 'content oblivious' investigation methods, like user reporting, are 'considered more useful than monitoring the contents of users' communications when it comes to detecting nearly every kind of online abuse.' That remains just as true in Florida today.

"Law enforcement can and does already conduct plenty of investigations involving encrypted messages. And even with end-to-end encryption, law enforcement can potentially access the contents of most messages on the sender or receiver's devices, particularly when they have access to the physical device. The bill also includes measures prohibiting minors from accessing any" - get this, Leo - "any sort of ephemeral messaging features" - they're taking that away, features away from minors - "like view once messages or disappearing messages. But even with those features, users can still report messages or save them. Targeting specific features does nothing to protect the security of minors, but it would potentially harm the privacy of everyone.

"SB 868/HB 743 radically expands the scope of Florida's social media law HB 3, which passed last year and itself has not yet been fully implemented as it currently faces lawsuits challenging its constitutionality. The state was immediately sued after the law's passage, with challengers arguing the law is an unconstitutional restriction of protected free speech."

Leo: Yeah.

Steve: "That lawsuit is ongoing, and it should be a warning sign. Florida should stop coming up with bad ideas that cannot be implemented. Weakening encryption to the point of being useless is not an option. Minors, as well as those around them, deserve the right to speak privately without law enforcement listening in. Florida lawmakers must reject this bill. Instead of playing politics with kids' privacy, they should focus on real, workable protections, like improving consumer privacy laws to protect young people and adults alike, and improving digital literacy in schools."

So exactly right, EFF. And I sure hope that the U.S. Supreme Court, Leo, doesn't mind working and being busy...

Leo: It's been kind of busy lately.

Steve: Oh, my god. I don't recall any time during my life when more important and fundamental issues surrounding the shape of our collective future are being pushed up our legislative hierarchy for their, the Supreme Court's, final examination, hopefully some useful discussion and judgment. And I sure hope they get these things right. They really are important.

Leo: Yeah.

Steve: One last piece, and then we'll take our final break, and then we're going to talk about Windows Sandbox. I found an interesting piece of reporting which I have - it was in Dutch, which I had Firefox translate. After examining more than 500 ransomware incidents occurring between 2019 and 2023, a Dutch researcher found that ransomware victims who are insured against the cost of cybercrime incidents pay, on average, 2.8 times larger ransoms than those who are uninsured.

Leo: Because it costs them nothing.

Steve: Uh-huh. And the bad guys, it turns out, know this. They make a concerted effort to research and determine the cyber-insurance status...

Leo: Oh, my.

Steve: Uh-huh, of all potential targets. The researcher wrote: "As soon as they have gained access to a system, they actively look for documents with names such as 'insurance' or 'policy.' This additional information gives cybercriminals a better bargaining position, leading to higher ransom payments." The research also found that companies with a well-designed backup system pay - get this, everybody. Companies with a well-designed backup system, and Leo, you're going to want to hold onto this factoid for your backup sponsors - pay 27 times less often.

Leo: Oh. Whew. That was a relief. Yes. Of course they do, because they've got a backup.

Steve: They don't need to pay.

Leo: Yes.

Steve: Twenty-seven times less often in the event of cyber attacks. The researcher wrote: "Cybercriminals who are in a victim's network consciously look for backups and remove them. Just having backups is not enough. It is important to have backups that cannot be adjusted by unauthorized persons in your network. Offline backups are the easiest solution for that," he said, "but I've also seen cloud solutions coming by." Meaning, you know, being a problem. So and the researcher also found that most companies have no choice other than to pay.

The researcher wrote: "In only around five out of 100 cases" - and he looked at 500, so 5% - "in which payments are made, victims do have the opportunity to recover in a different way than to pay, but choose to pay anyway, for example, to recover faster or to prevent reputational damage. In other words, in only five out of the 100 cases, they will voluntarily pay the ransom for the ancillary benefits aside from the ability to continue being a viable business. In the other 95 out of 100 cases," he wrote, "there is no other option to recover. In those cases, their entire IT infrastructure is broken and no longer recoverable, making paying a ransom the only option to prevent their bankruptcy."

Leo: Wow.

Steve: So I suppose it's not really surprising that 95% of ransomware victims do not have a sufficiently comprehensive or attack-proof backup system in place. So they really do have no other choice than to give the extortionists whatever it is they demand. It's either that, go out of business, or start again from scratch. And we know from our own years of looking at this that the bad guys will also actively look for and work to eliminate any backup systems and servers that they can find. They'll crash those and then, you know, wipe them, and then exfiltrate the data and encrypt, you know, everything that remains. They're also aware of that 95/5 rule, that 95 out of 100 cases the company has no choice but to pay. And they very much want their victims to have no other recourse than to pay them. Really interesting data.

Leo: Yeah, yeah, very interesting, yeah.

Steve: So, okay. Often ignored or unknown to most users of Windows 10 and 11, but probably of tremendous value and interest to the followers of this podcast, is that built right into every Win10 and Win11 64-bit Pro, Enterprise, and Education operating system - Home is the only edition that doesn't have it - is a ready-to-use, extremely robust, virtual machine-based full security sandbox inside of which Windows users can perform any experiments they may wish where everything they or their experiments do will deliberately be "sandboxed" from the enclosing host PC and will therefore be unable to affect or in any way damage the hosting PC.

And what's surprising is that this quite valuable security feature has been right there, available and in front of us since 2018 with the release of Windows 10 version 1903. And because it's not enabled or installed by default, mostly we're unaware of it. But oh, wait till you hear about the technology. I mean, as I said at the top of the show, rarely am I impressed, I guess is the best word, with what Microsoft does. This thing, I am infatuated with it.

Microsoft describes this sandbox built into Windows 10 and 11 by writing: "Windows Sandbox is a lightweight, isolated desktop environment designed for safely running applications. It is ideal for testing, debugging, exploring unknown files, and experimenting with tools. Applications installed within the sandbox remain isolated from the host machine using hypervisor-based virtualization. As a disposable virtual machine, Windows Sandbox provides quick launch times and a lower memory footprint compared to VMs." And wait till you understand why that's not just marketing BS.

For key features, Microsoft highlights: "Part of Windows: Everything required for this feature is included in supported Windows editions like Pro, Enterprise, and Education. There's no need to maintain a separate VM installation. Disposable: Nothing persists on the device. Everything is discarded when the user closes the application. Pristine: Every time Windows Sandbox runs, it's as clean as a brand-new installation of Windows.

Secure: Uses hardware-based virtualization for kernel isolation. It relies on the Microsoft hypervisor to run a separate kernel that isolates Windows Sandbox from the host. And Efficient: Takes a few seconds to launch, supports virtual GPU, and has smart memory management that optimizes its memory footprint."

So this is clearly a win for anyone who might have any occasion to need a quick, safe, disposable instance of Windows. Because that's what you get. You are booting a brand new Windows that's built right in. Windows on Windows. You know, perhaps you'd like to install something to see what it looks like, but it's a big lumbering thing that's likely to change your icons and create file associations and reconfigure and whack a big portion of your finely-tuned desktop, so you haven't installed it because of the hassle of probably later uninstalling it, you know, and then maybe recovering your machine from, you know, from everything it did to it. It's just not worth the trouble of just, like, satisfying your curiosity.

With Windows Sandbox there's nothing to uninstall. Just close the instance of Windows running in Windows, and that monstrosity will be gone like it never existed the next time you use the sandbox. Or perhaps there's a sketchy program you found, you know, somewhere on the Internet which you'd really like to run, but haven't dared to on any machine that might be hurt by it. Or perhaps you need to poke into some particularly dark corners of the Internet, don't want anything to poke you back, and don't want to leave a trace, want to leave absolutely no trace of ever having done that. It turns out that every non-Home edition of Windows has this capability built right in and ready to do all those things.

And what's so extra cool about this is that the Windows Sandbox is able to be far more efficient than a traditional full virtual machine setup. It's able to adjust its memory usage according to the demand, and it doesn't require an entire second installation of Windows since it's able to reuse many of the host's read-only operating system files. It is quite a slick solution.

Okay. So first, where is it, and how do you obtain the use of this little forgotten gem? On the desktop, enter the search bar, search for "Windows Features." It's actually "Turn on and off Windows features," so you can probably type in "turn on and off," and that would get it.

Leo: Oh, that's a classic control panel. That's still around?

Steve: Yup.

Leo: Loved that, yeah.

Steve: Still there. That will bring up the, you know, that Windows Features panel. And you'll see most of the things there are turned off. They're mostly things that are kind of optional, like most people don't need to run IIS, the Microsoft Internet Server. Or there are some enterprise-y things. Most of the checkboxes are off. Scroll down near the bottom, and I think it's like fourth from the bottom, you will see a checkbox you've probably never noticed before, or didn't think about it, labeled "Windows Sandbox." You just check that box to turn the feature on, then click "OK" to confirm your choices. Windows will spend a minute or two unpacking its bags, and will then tell you that you need to reboot to finish things up.

Now, there's a chance that you may find that feature greyed out and unselectable. If you hover your mouse over that, Windows will probably inform you that Windows Sandbox cannot be installed because the processor does not have the required virtualization capabilities.

Leo: Ah, ah.

Steve: If that's the case, you may be able to remedy that by rebooting, getting into your machine's firmware BIOS, or UEFI, and enabling the various processor virtualization features that are needed. Like I'm a fan of VirtualBox, and sometimes when I'm setting it up on a new machine it won't run because I need to go into the BIOS and turn on, you know, VT-X in order to enable the virtualization features in the processor which need to be turned on by the firmware at boot time. So the other possibility is that you might be in a VM trying to run Windows in a VM. So if you scroll down to the bottom, Leo, you will find...

Leo: Oh, I see, it's greyed out. But that's because I'm running in virtualization.

Steve: Ah, yes.

Leo: You couldn't have a sandbox in a virtual environment, probably.

Steve: Actually you are able to.

Leo: Ah.

Steve: You are able to turn - but you need to enable virtualization within virtualization, which is a feature of the virtualizing systems.

Leo: Ah. I will go examine that. That's great.

Steve: You probably are able to do that. So after you reboot, if you scroll down in the main menu, down into the "W's," you will find listed all there by itself, I mean, like along with Windows applications or Windows administration and Windows systems down there, and in fact it's right above Windows system, which is a folder that expands, is Windows Sandbox all by itself. Click it, and you will shortly be presented with something you may dimly recall, which is your original Windows system before it was first touched.

Leo: Ooh.

Steve: It is completely clean.

Leo: Nice.

Steve: Nothing installed. Yes. You'll see it looks like a standard Windows window named "Windows Sandbox." It's got the minimize, maximize, and close icons in the upper right, as Windows apps do. I noticed that resizing the virtual machine window was as smooth as anything I've ever seen.

Leo: Sure. There's nothing else running in the background.

Steve: Yes. Well, but, I mean, even the host system doesn't seem to have any problem hosting what is another running Windows boot. I mean, it booted Windows.

Leo: It sounds like it's a little bit like Docker. Is it like Docker? Do you know how Docker works?

Steve: I don't know how Docker works.

Leo: So one of the nice features of Docker is it doesn't install an entire operating system. It runs another operating system, but it uses the operating system resources already there.

Steve: That's exactly what this does.

Leo: Yeah, it's clever.

Steve: I will be explaining that in a second. Yes, that is exactly what it does. So I also noticed that if I maximized the window, it just became my desktop. It completely took over and covered up the underlying hosting desktop, and it showed the Remote Desktop Connection bar at the top center. So Remote Desktop is the way the virtual machine's desktop is being presented to the user.

The sandbox has a C: drive with about 3GB shown as being in use - although it actually doesn't take up 3GB, we'll get to that in a second - and plenty of empty space. Internet access by default with a generic LAN adapter is present, so you have Internet access from within the sandbox. It's got the IP address of 172.17.*.*, whatever, an RFC, what is it, 1913 private network that is set up, and it has a single user account named "WDAGUtilityAccount," where WDAG stands for Windows Defender Application Guard. However, Microsoft notes that Windows Defender does not actually run inside the Windows Desktop. Again, they're trying to keep it fast and lightweight. And as many people know, Windows Defender can sometimes start up and slow things down for a while, while it's scanning through everything.

Anyway, Microsoft really appears to have done a nice job of this. I was curious to see what would happen if I attempted to launch a second instance of the Sandbox, and I was greeted with a dialog from Windows Sandbox that said: "Only one running instance of Windows Sandbox is allowed." So, okay. I closed that. And then out of curiosity, I tried clicking the upper-right close "X" and was told: "Are you sure you want to close Windows Sandbox? Once Windows Sandbox is closed, all of its content will be discarded and permanently lost." Which of course is exactly what we want.

And the second time the Windows Sandbox is launched, its desktop pops right up, though that's somewhat misleading since Windows is not actually ready, and it does still need a bit more time to get itself actually booted. You know, as the old timers among us will recall, at one point Microsoft was receiving so much flack over how long Windows was taking to boot that they deliberately engineered it to display its desktop at the earliest possible moment after, like, turning the machine on and getting it to start booting, which was well before it was actually able to do anything. I always thought all that ingenuity would have been better spent actually making it boot faster, but no one asked me.

Anyway, before we dig under the covers to take a closer look at the technology behind this, let's look at some more of the surface details. Windows Sandbox is also available on ARM64 from Windows 11, version 22H2 on. So you can get it for ARM and Intel platforms both, or AMD64 of course also. And starting with Win11 24H2, the inbox store apps like Calculator, Photos, Notepad, and Terminal are not available inside Windows Sandbox. They said that the ability to use these apps is going to be coming soon.

A so-called "vGPU," a virtualized GPU, is enabled on non-ARM64 devices. As I noted, networking is enabled using the Windows Hyper-V default switch. Since this could potentially expose untrusted applications to the user's internal network, it is possible to launch a Sandbox with networking disabled or to disable it, you know, after the fact through the use of a custom .wsb file, as in "Windows Sand Box" configuration file. Audio input is enabled with the sandbox by default having access to the host's microphone input. But video is not by default. The sandbox does not share the host's video, or the host does not share its video with the sandbox.

Printer redirection is also disabled, with the sandbox not sharing printers with the host. But keyboard redirection, I'm sorry, clipboard redirection - of course keyboard is. But clipboard redirection is enabled by default so that the host's clipboard is shared with the sandbox, allowing for the cutting and pasting of text and filenames, you know, back and forth, which is just a convenience.

It's also possible to change all of those defaults and many other aspects of Sandbox's configuration. Windows Sandbox supports, as I mentioned, that WSB, which is a simple XML format configuration file, which provide a minimal set of customization parameters for the Sandbox. This feature can be used with Windows 10 build 18342 and later, or Windows 11. So that wasn't quite in that earlier 1903, but 1842 or later. Windows Sandbox configuration files are formatted, as I mentioned, as XML, and are associated with the .wsb file extension.

A configuration file, that little .wsb, enables the user to control a number of aspects of the sandbox. That virtualized GPU can be disabled to cause the sandbox to use Windows Advanced Rasterization Platform, known as WARP. Networking can be disabled. Mapped folders can be defined to allow the sandbox to see some controlled aspects of the host's file system, if you like. A custom logon command can be executed when the sandbox starts. The audio and video sharing defaults can be changed to either allow or disallow video and audio.

Remote Desktop Protocol's "Protected client" mode, which is an elevated level of security, can be engaged to place that increased security settings on the Remote Desktop Protocol session which is used to access the sandbox. Printers can be shared, the clipboard sharing can be disabled, and the total amount of memory assigned to the Sandbox can be changed from its default of a hopeful 4GB, although it will use less if less is available.

Okay. So now I want to turn the clock back to December at the end of 2018 and look at what Microsoft shared about this terrifically useful innovation back then. The Windows OS Platform Blog posted, under the simple title "Windows Sandbox," they said: "Windows Sandbox is a new lightweight desktop environment tailored for safely running

applications in isolation. How many times have you downloaded an executable file, but were afraid to run it? Have you ever been in a situation which required a clean installation of Windows, but didn't want to set up a virtual machine?" Or for that matter even another real machine.

They wrote: "At Microsoft we regularly encounter these situations, so we developed Windows Sandbox: an isolated, temporary desktop environment where you can run untrusted software without the fear of lasting impact to your PC. Any software installed in Windows Sandbox stays only in the sandbox and cannot affect your host. Once Windows Sandbox is closed, all the software with all its files and state are permanently deleted. Since this is the Windows Kernel Internals blog, let's go under the hood. Windows Sandbox builds on the technologies used within Windows Containers." Which, Leo, is presumably like Docker, as you said.

Leo: Yeah, containers, yeah, yeah.

Steve: "Windows containers were designed to run in the cloud. We took that technology, added integration with Windows 10, and built features that make it more suitable to run on devices and laptops without requiring the full power of Windows Server. Some of the key enhancements we have made include a dynamically generated Image. At its core, Windows Sandbox is a lightweight virtual machine, so it needs an operating system image to boot. One of the key enhancements we've made for Windows Sandbox is the ability to use a copy of the Windows 10 installed on your computer, instead of downloading a new VHD image as you would have to go through with an ordinary virtual machine.

"We want to always present a clean environment, but the challenge is that some operating system files can change. Our solution is to construct what we refer to as a 'dynamic base image,' an operating system image that has clean copies of files that can change, but links to files that cannot change that are in the Windows image that already exists on the host. Again, no duplication of resources. The majority of the files are links, immutable files, and that's why it has such a small size of around 100MB for a full operating system. We call this instance the 'base image' for Windows Sandbox. When Windows Sandbox is not installed, we keep the dynamic base image in a compressed package of around 25MB. When installed" - so that's what happens when you click the you want to enable Windows Sandbox in the Turn Windows Features On and Off menu. "When installed, the dynamic base package it occupies is expanded to 100MB of disk space."

Okay. "So what about memory? Memory management is another area where we've integrated with the Windows Kernel. Microsoft's hypervisor allows a single physical machine to be carved up into multiple virtual machines which share the same physical hardware." Okay, that's standard VM technology; right? But while that approach works well for traditional server workloads, it isn't as well suited to running devices with more limited resources. We designed Windows Sandbox in such a way that the host can reclaim memory from the sandbox if needed.

"Additionally, since Windows Sandbox is actually running the same operating system image as the host, we allow Windows Sandbox to use the same physical memory pages as the host for operating system binaries via a technology we refer to as 'direct map.' In other words, the same executable pages of NTDLL, the kernel, are mapped into the sandbox as on the host. We take care to ensure this is done in a secure manner, and no secrets are shared."

Okay. So I imagine everybody can detect how utterly infatuated I am with this technology. It is genius. They're reusing all of the Windows OS files. They're reusing all of the Windows kernel's memory that's been loaded with static code. So any entirely separate and clean instance of Windows only requires around 100MB of storage, which is essentially a file system full of pointers into the host's file system. And rather than needing to create another virtual machine with its own allocation of 4GB or more of RAM, it also takes almost no RAM to run because it's able to map most of the host's actual physical RAM into its own virtual image. It is a win.

And there's more. They write: "With ordinary virtual machines, Microsoft's hypervisor controls the scheduling of the virtual processors running in the VMs." But I'll note they don't control the scheduling within the VMs, which is the key. They wrote: "However, for Windows Sandbox we use a new technology called 'integrated scheduler,' which allows the host to decide when the sandbox runs.

"For Windows Sandbox we employ a unique scheduling policy that allows the virtual processors of the sandbox to be scheduled in the same way as threads would be scheduled for a process. High-priority tasks on the host can preempt less important work in the sandbox. The benefit of using the integrated scheduler is that the host manages Windows Sandbox as a process rather than a virtual machine, which results in a much more responsive host, similar to Linux KVM. The whole goal here is to treat the Sandbox like an app, but with the security guarantees of a Virtual Machine."

And that's the genius of this. It really is running an entirely separate instance of Windows like an app on the underlying host OS. When you click it and launch it from the Start Menu, it's like you are just running an app, but that app happens to be a completely clean instance of Windows in which nothing has ever been done or installed, ready for you to play with. And remember how I mentioned that when I launched the Sandbox a second time it seemed to snap right up? This blog explains why I experienced that, too. It wasn't just my imagination, or my infatuation.

They wrote: "As stated above, Windows Sandbox uses Microsoft's hypervisor. We're essentially running another copy of Windows which needs to be booted, and this can take some time. So rather than paying the full cost of booting the sandbox operating system every time we start Windows Sandbox, we use two other technologies: 'snapshot' and 'clone.'

"Snapshot allows us to boot the sandbox environment once and preserve the memory, CPU, and device state to disk. Then we can restore the sandbox environment from disk, loading it directly into the device memory, rather than booting it, when we need a new instance of Windows Sandbox. This significantly improves the start time of Windows Sandbox." Essentially, once Windows finishes booting the first time, they snapshot all of the work that was done to get it booted, and save that, too. So that when you relaunch Windows, it comes up, and then it restores the virtual machine state from which that snapshot was made.

And graphics virtualization. They said: "Hardware-accelerated rendering is key to a smooth and responsive user experience, especially for graphics-intense or media-heavy use cases. However, virtual machines are isolated from their hosts and unable to access advanced devices like GPUs. The role of graphics virtualization technologies, therefore, is to bridge this gap and provide hardware acceleration in virtualized environments. More recently, Microsoft has worked with our graphics ecosystem partners to integrate modern graphics virtualization capabilities directly into DirectX and WDDM, the driver model used by device drivers on Windows. Graphics components in the Sandbox, which have been enlightened" - I like that, which have been enlightened - "to support virtualization, coordinate across the VM boundary with the host to execute graphics workloads.

"The host allocates and schedules graphics resources among apps in the VM alongside the apps running natively on the host. So essentially the boundaries have been softened as much as they possibly could be so that there is really no difference between apps running in the Windows Sandbox as apps running on the host desktop."

They said: "This enables the Windows Sandbox VM to benefit from hardware accelerated rendering, with Windows dynamically allocating graphics resources where they're needed across the host and guest. The result is improved performance and responsiveness for apps running in Windows Sandbox, as well as improved battery life for graphics-heavy use cases. To take advantage of these benefits, you'll need a system with a compatible GPU and graphics drivers (WDDM 2.5 or newer)." Remember this was written in 2018, so we probably all have that. "Incompatible systems will render apps in Windows Sandbox with Microsoft's CPU-based rendering technology.

"And finally, battery pass-through: Windows Sandbox is also aware of the host's battery state, which allows it to optimize for power consumption. This is critical for a technology that will be used on laptops, where not wasting battery is important to the user."

So I've been spending a lot of time recently using Virtual Machines. The DNS Benchmark that I'm currently working on needs to run under Windows 7, 8, 10, and 11, and those four operating systems span enough time that their behavior is all slightly different from one another. So I am routinely launching and running different OSes on different platforms. When I originally built my main old Windows 7 machine, I expected virtualization to be a thing that I would want to have access to, so I deliberately gave it a whopping 128GB of main system memory. This was specifically so that I could fire up separate Windows virtual machines that would each need large chunks of RAM dedicated for their own use. And my Windows 10 machine has 32GB, which was the most that that Intel NUC could handle at the time.

My point is that I've become quite accustomed to the feeling of virtual machines running on my desktop, and I have never experienced as seamless and smooth an operation of a Windows OS in an OS as is provided by this built-in Windows Sandbox. I really believe Microsoft has outdone themselves on this one. They've been very clever, and they've done everything right. They've essentially figured out how to run an entire separate instance of Windows as an application, and even the applications in that as applications for the host on top of Windows. It's fast and lightweight and does not burn up disk space or RAM.

Anyway, toward the end of today's show notes I have a collection of links to additional resources to help anyone get the most out of their built-in Windows Sandbox, including all the documentation about configuring and tweaking its operation, RAM, cross-host sharing, resources, shared folders and everything else. Anyway, there are so many really compelling use cases for this slick technology that I wanted to make sure all the listeners of this podcast who use Windows as their primary desktop knew that this little gem was hidden right there, I mean, just waiting to come out to play.

I am, as I mentioned at the top of the show, I'm finally somewhat jealous of Windows 10 - actually I'm sitting in front of it right now, but I don't have it on my Windows 7 system, where I've not been in any hurry to upgrade the Windows 7 machine because everything works just fine. But now I'm thinking maybe I'm going to take a big, one final, you know, system image snapshot and then see if Windows 10 is able to upgrade from my old Win7 machine. I dread the downtime required to set up a new Windows 10 machine from scratch and reinstall everything and configure it all, I mean, that's just days of work. But Windows Sandbox has been implemented so beautifully that it's something I would love to have on that other desktop platform. I mean, Leo, it's just, I mean, it's just a - they did a beautiful job.

Leo: Nice.

Steve: And again, containment. It is a security sandbox, so I would expect it would be of tremendous interest to our listeners. And they already have it. They just most of them, like me, I'd forgotten about it.

Leo: How funny. Yeah.

Steve: I remember it once upon a time, but I'd completely forgotten it.

Leo: Right, right. It's been there all along.

Steve: Yeah.

Leo: Hiding in plain sight.

Steve: I mean, many instances where I've downloaded something sketchy. Remember when I was doing all the work on SpinRite 6.1, and I needed networking drivers for long obsoleted network adapters. And I had to, like, download things from sketchy sites in order to get the DOS drivers. And I was like, eeeeeeeeegh. Well, I could have unpacked them in the sandbox and then just taken the files themselves safely and not worried that the zip file might have been compromised with some sort of other goo.

Leo: Super cool.

Steve: Really, really neat.

Leo: Super cool. Yeah, containers are a good thing. I think it's a very exciting area right now in the community.

Steve: It is. The idea of reusing the static footprint of an operating system and its static files, it makes so much sense.

Leo: Sure. Why have duplicates in RAM; you know?

Steve: And I think next week I'm going to share how malware has decided to move into Windows Sandbox.

Leo: Yeah, somebody was saying it's just a matter of time before we have a story about sandbox escapes. But that's a - we'll save that for future.

Steve: It's actually - it's actually not an escape because the isolation is extremely good, although not to say that there might not be some.

Leo: Right.

Steve: But it turns out malware is using the Windows Sandbox to hide.

Leo: Oh, that's a good idea. That makes sense. All right, we'll talk about it next week.