The Little Red Wagon

Description: This week we examine the most recent flaw found in Intel's processors and what it means. We look at the continually moving target that is Windows 10. We consider the Free Software Foundation's suggestion that Microsoft open source Windows 7 and the fact that last month's was apparently NOT the last update of Windows 7 for all non-ESU users. We look at the evolution of exploitation of the Remote Desktop Gateway flaw, Google's record breaking vulnerability bounty payouts, the return of Roskomnadzor, the size of fines, the question of who owns our biometrics, an update on Avast/AVG spying, the future of third-party AV, a major milestone for the WireGuard VPN, and the wonderful Little Red Wagon hack of the decade which titled this podcast.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-752.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-752-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. It's one of those great potpourri episodes with lots to talk about, including the strangest Windows 7 bug. Microsoft says this time they are going to fix it. We'll also talk about some scary hacks, including the latest Intel problem. And then that Little Red Wagon that could. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 752, recorded February 4th, 2020: The Little Red Wagon.

It is time for Security Now! with our friend and chief, Steve Gibson, the man in charge of the Gibson Research Corporation. I've always meant to ask you. What is it you're researching, Steve?

Steve Gibson: Well, the story behind that is that I really liked Gibson Labs as the name. And that was my previous company's name, Gibson Labs. That's where I did the light pen.

Leo: That must be where I got the idea for Tech Guy Labs, from that. Because that's the website for the radio show.

Steve: Maybe. I just sort of liked that. But I sort of sold it to Atari along with the light pen.

Leo: Boom. Whoops.

Steve: And when I wanted then to do something again, my attorneys at the time said, you know, technically...

Leo: You sort of sold that.

Steve: You still have it. Well, because then Atari backed out because the home computer business collapsed. So then I took it to Koala, and they made the light pen for a while. And then...

Leo: I remember them.

Steve: And then it kind of stayed there, but it wasn't clear whether they just got the pen or everything. And so my attorneys said, "Yeah, you need another name." And I said, "I don't want another name." And they said, "You need another name."

Leo: You've got to have one.

Steve: I said, "Okay, Gibson Research."

Leo: There you go.

Steve: I'll research something.

Leo: Something. Anything.

Steve: And they said, "Okay, that's what we'll do." So that's, you know.

Leo: That's really funny.

Steve: I could probably get it back now, but I don't care. Now it's GRC. And I mean, like, well, and it's funny, too, because I remember when we were trying to get a domain, I said to Millard Ellingsworth III...

Leo: What?

Steve: Who was, yeah, he was one of my guys. I said, "Millard." And he was sort of that way anyway. I said, "We need to have a domain name, apparently." And this was, like, six months after Microsoft, Bill, had decided the same thing. So of course they already had their name. Anyway, so Millard came back, he said, "Well, we can't get Gibson." And I said, "Why?" He says, "Guitars." I said, "Oh."

Leo: Oh, yeah. Oh, yeah.

Steve: Or maybe it was refrigerators. I don't remember. It was something.

Leo: There is a guitar company called Gibson.

Steve: Yeah. I do know that. And he said, "But I got GRC." And I said, "Oh. I think I like that." And so...

Leo: It's short.

Steve: Yeah, exactly. And then, like, sometime later, because back then we were using something called cc:Mail, where we had sort of a network, and we had a dialup that would phone into something and retrieve our email.

Leo: I remember that.

Steve: And retrieve our email every so often.

Leo: Yeah, yeah, yeah.

Steve: You know, I mean, it was the dark ages still. And it was like a big deal when I wired the office for coax in order to have 10Base2 interconnectivity, and we could all print to a common printer. And we had a laser printer; and that was, like, whoa, look at that, doesn't make any noise. And so eventually it was like time to actually get on the Internet. And so there was a local company down here, and so I called them up. And this gal answered, and I said, "Yeah, I need to hook up to the Internet." So she says, "Okay. Do you have a domain?" I said, "Yes, I do." And she said, "What is it?" And I said, "GRC.com." And even back then she said, "Oh, three letters."

Leo: Three letters, whoo.

Steve: And I get offers constantly. The most recent one was an offer of $50,000 for just a blind offer for...

Leo: What?

Steve: Yeah, they just want GRC.com because it's three letters.

Leo: That's impressive. Wow.

Steve: Yeah, so, you know, somehow, after I'm dead, someone, my estate will liquidate GRC.com and, you know...

Leo: Get profit.

Steve: By a little bit. Today's Episode 752 - now, because you're clued into the world, Leo, you knew what I meant by "The Little Red Wagon."

Leo: That's the title.

Steve: And that's the title of today's podcast is "The Little Red Wagon," just because it is the coolest hack that - anyway, I can't wait to share it with our listeners. That will be - we didn't have like a big earth-shaking killer topic. Lots of interesting stuff happened, which is what today's podcast will be about. But it will conclude with what is probably the coolest hack I have seen of the decade. And I'm not sure whether the decade actually starts next year or this year. There seems to be some concern about that. But it's just so fun. So "The Little Red Wagon" is the title for today's podcast.

Leo: I added "Little" because it just said "The Red Wagon." So I put "Little" in there.

Steve: Oh, and that's interesting because - maybe it wouldn't fit. Because I'm looking at my show notes, and it says "The Little Red Wagon."

Leo: Oh, yeah. You got it right, yeah. Somebody just left out "little," that's all.

Steve: Ah, okay. And believe it or not, the Picture of the Week is the way it came. It just happens to use my name, but that's a fluke.

Leo: All right, Steve. This is the customized Picture of the Week.

Steve: Okay. So, yeah, it just happened to involve a character named Steve. We've got two guys walking down the sidewalk. And they've just passed a store whose awning says "Clothes for Steve." And then they're crossing in front of the store "Furniture That Would Look Good in Steve's House." And then down toward the end of the whole cartoon is "Steve's Favorite Ice Cream Flavors."

Leo: Okay.

Steve: And so the caption underneath says - the other guy's apparently talking to Steve, says "Maybe you should disable your cookies, Steve."

Leo: I love it.

Steve: Yes, you are being tracked.

Leo: And really the web is like that; isn't it. Hi, Steve.

Steve: It is very much like that.

Leo: Yeah, everything's for Steve.

Steve: Okay. So our first little tidbit, we've got lots of little tidbits, was the follow-up on the news that you broke at the beginning of last week's podcast, L1D eviction sampling, which the news had just happened. And you presented this, and correctly, as a breaking news item. Another flaw in Intel chips had just surfaced called L1D Eviction Sampling. And we guessed correctly that L1 would refer to the Level 1 cache. And since it had the term "eviction" in it, which is the term used by caching, in fact, I heard you use it on MacBreak Weekly. You were talking about CacheFly and evictions, eviction of cached data. So of course it applies there, too. Eviction is the term used when new room needs to be made in the cache for something more recent.

And so typically LRU caching, Least Recently Used caching, knows what the oldest thing is which is just sort of statistically less likely to be reused. And so it gets evicted from the cache. Anyway, so we were right about our guess. We now have the fancy name, or the popular name. It's called CacheOut. And of course it has a website, CacheOutAttack.com. And the logo, it has to have a logo these days, and so it's got a slot machine as the CacheOut Attack logo.

The page's introduction explains - this is the "Leaking Data on Intel CPUs via Cache Evictions." And they said, the researchers said: "We present CacheOut, a new speculative execution attack that's capable of leaking data from Intel CPUs across many security boundaries. We show that, despite Intel's attempts to address previous generations of speculative execution attacks, CPUs are still vulnerable, allowing attackers to exploit these vulnerabilities to leak sensitive data."

Then they said: "However, unlike previous MDS" - and now we know, we've all been tuned up on this, this is the acronym for Microarchitectural Data Sampling, so MDS, Microarchitectural Data Sampling issues. "We show in our work how an attacker can exploit the CPU's caching mechanisms to select what data to leak, as opposed to waiting for the data to be available. Finally, we empirically demonstrate" - and do they - "that CacheOut can violate nearly every hardware-based security domain, leaking data from the OS kernel, co-resident virtual machines, and even SGX enclaves."

So these are researchers at the universities of Adelaide and Michigan. They showed in their paper - and I'll share five bullet points because these are a little bracing - the effectiveness of CacheOut in violating process isolation by recovering AES keys and plaintexts from an OpenSSL-based victim. Second bullet, practical exploits for completely de-randomizing Linux's kernel ASLR - right, the Address Space Layout Randomization, which is like a preamble for then exploiting return-oriented programming attacks - they said, and for recovering secret stack canaries from the Linux kernel. And that again, if you can recover the stack canary, then you're able to avoid tripping Linux's stack overflow checks as again, in order to leverage a powerful attack.

Three, how CacheOut effectively violates the isolation between two virtual machines running on the same physical core. Four, how CacheOut could also be used to breach the confidentiality SGX guarantees by reading out the contents of a secure enclave. And,

finally, how some of the latest Meltdown-resistant Intel CPUs are still vulnerable, despite all of the most recent patches and mitigations.

And as I was pulling the notes together, I then wrote at this point: "Intel, who really used to enjoy their original job of printing money, once again responded to this latest annoyance." Because, gee, it's just not so easy to make these checks anymore. We used to just go out in the back, and we'd get a wheelbarrow full of sand, which is of course silicon, and melt it down and purify it and then just charge a thousand dollars for these little tiny chips of silicon. That's not quite so easy anymore.

So basically they acknowledged, yes, they have now the security software guidance. And it's like, once upon a time, firmware, I mean, microcode updates were like not a thing. Now, yeah, get your latest update here.

So they said: "A speculative execution side channel variant known as L1D Eviction Sampling may allow" - we know that everybody, you know, Microsoft phrases their things the same way, as opposed to "has been shown to allow." No, "may allow the data value of some modified cache lines in the L1 data cache to be inferred under a specific set of complex conditions." Yeah, well, okay, fine. That's not much of a mitigation because bad guys, as we know, are willing to work hard to infer the data, the value of the data that's just been evicted from the cache.

Intel said: "On some processors under certain microarchitectural conditions" - which is to say ours - "data from the most recently evicted modified L1 data cache line may be propagated into an unused, invalid L1D fill buffer." Okay. Basically they're going to dazzle us with the details. "On processors affected by Microarchitectural Data Sampling or Transactional Asynchronous Abort (TAA), data from an L1D fill buffer may be inferred using one of these data sampling side channel methods. By combining these two behaviors together, it may be possible for a malicious actor to" - and again, has been shown to be possible - "to infer data values from modified cache lines that were previously evicted from the L1 data cache." At this point everyone starts to snore.

They said: "This is called L1D eviction sampling. Malicious software may be able to use L1D eviction sampling to infer modified cache line data written by previously run software, or modified cache line data written by software running on a sibling hyperthread on the same physical core." And note here the term "modified" in "modified cache line data" refers to the fact that the data being evicted has been modified, and that cache is "write back" rather than "write through." So that modified line needs to be written back out to at least a lower level cache, maybe all the way back out to external DRAM. That's what takes the time. If the data had not been modified, then it could just be overwritten, which would take no time.

Anyway, they go on with similar, like, you know, I won't drag our listeners through it, but mumbo jumbo of this sort. So the CacheOut page has an extensive FAQ for anyone who's interested. But the takeaway for most of us, if not all of us, is the same as it's been for the past two years. Yes, it's scary, and it sounds bad. But even though it's real, it's another difficult-to-exploit fringe theoretical academic problem. No known real-world attacks have ever been detected in the wild.

On the other hand, they would not likely be because it doesn't leave any obvious footprints. It just allows one evil process sharing a core, because that's where the L1 cache is, so it allows - you have to have core sharing; two, if it wanted to find something out about what has recently been done on that core. And Intel has released another round of microcode patches for this, which reduces processor performance, of course, because that's the way that you have to deal with speculation. The reason all of these performance improvements have been put into the microarchitecture is to creep the

performance forward. Linux can incorporate these newly released microcode patches into its boot, and Windows 10 will be getting them. In fact, has them.

So basically, interesting, and I think important for us to just sort of address for the sake of understanding the territory that we're now operating in. And it will be nice when we start having high-performance Intel cores that don't have all of these problems, if we can get the best of both worlds, and it's not yet clear we can. But, for example, the way to resolve this would be to turn off hyperthreading. Except hyperthreading is an inexpensive means of keeping a core busy, inexpensive inasmuch as it allows two threads to do a very quick context switch, to jump between threads, and where you're only storing a little bit about what each thread's state is. And so Intel doesn't want to give it up.

And users who were super concerned about this could turn off hyperthreading, and normally that's still an option in the BIOS. Just I don't know why it's there, but it's kind of handy that it's there now. Not that anyone's really ever had any problem with it. So the researchers noted that it's unlikely that AV products would detect or block CacheOut because it would be really difficult to see it happening. And it's very unlikely that anybody has exploited it.

So it is not in AMD processors. Apparently there were a couple brands, I think it might have been - I don't remember now. I think I saw IBM and one other processor sort of was doing the same thing. They had not looked at those. So we don't know. But certainly anyone who's making processors these days is looking at what's happening with all of the incoming fire that Intel is taking and thinking, ooh, let's fix that before one of those annoying academic researchers thinks, hmm, I wonder what about these processors. And speaking of microcode updates for Windows, which can be loaded at boot time, last Thursday Microsoft did release updated Intel microcode for Windows, pretty much all of them - 1909, 1903, 1809, 1803, 1709, 1703, and 1607. They also had one they just labeled Windows 10. And it's like, oh, were they calling it Windows 10 before 1607? I don't know.

Leo: All of the above.

Steve: Yeah. So the important part is this came out on Thursday, which was like, not a Tuesday; right? And it's all versions of Windows. Or, well, of Windows 10. So no Windows 7 fixes, significantly, but we didn't expect them any longer. Except, as we'll be seeing a little bit later in the show, there will be an upcoming change to Windows 7 due to something that they broke a couple weeks ago.

So in the show notes I've got - because it turns out that each of these microcode updates is different, depending upon the edition of Windows you're using. So if you're at 1903 or 1909, you need KB4497165. 1809 needs KB4494174. If you have 1803, you need KB4494451. And I'm going to explain in a minute why I'm going through this. If you're at 1709, you need KB4494452. 1703, KB4494453. 1607, KB4494175. And what they described as just Windows 10. And I went to the knowledge base article, and that's what it just says, Windows 10. It's like, oh. That's KB4494454.

So the point is, if you were concerned, or you had a version of Windows 10, or presumably Windows 10 Server - the server versions are similarly affected, and I would argue that only the server versions, which might be running untrusted code, would need to bother with these updates - then Microsoft is not giving them to you. You need to go to that knowledge base article and manually install the Windows Update from that article. And, boy, you look at any of those, and they cover everything, Leo. You ought to just google "KB4494454" and look at the list of processors. Because now they've added coverage for Denverton, Sandy Bridge, Valley View, and Whiskey Lake on top of, like,

everything else like all the way back that is being affected by this. So it's a mess. But I want to say I would never consider putting any of those on any of my machines, on any of my Windows 10 machines.

Leo: Because you're not in a shared environment.

Steve: Exactly. When or if Microsoft should ever decide it's necessary for me to have them, then I presume it'll be part of a monthly rollup, and I'll just get them, and I won't really be able to say no unless I try to fight back. There's really no reason for an individual to do that. We all cherish the performance of the machines we have. Caching and speculative execution are all about performance. So we don't want to turn that off.

And if you've got something in your personal machine which is trying to steal something, well, first of all, there's lots of easier ways to do it, if it's code running on your machine, than trying to leverage a very subtle timing flaw in L1 cache eviction. So you've got bigger problems than that, if you've got something in your machine trying to do that. So it's like, eh, again, if something were ever found to be doing it, then Microsoft would just roll that out to the rest of us. At this point I don't think it makes any sense.

And looking over those seven different knowledge base articles, each for one - or mostly for one. In the case of 1903 and 1909 they're sharing one. What the hell ever happened, Leo, to there being only one Windows 10 from now on? Wasn't that the big promise?

Leo: Well, they said the last version of Windows. I don't know.

Steve: Okay.

Leo: If they said there'd only be one.

Steve: That sounded like the one good thing that this new operating system had going for it. It was like, oh, thank god. They're finally going to give us Windows 10, and we'll be done. But as I've been gaining some experience with it, as far as I can see, what we now have is effectively separate and differing versions of Windows that don't change now every six years, the way they used to. They're changing every six months.

And I was just recently, last week, trying to solve a mystery with Windows 10, and getting it to work on an old version of Server Message Blocks, SMB v1. And what was really interesting was I encountered exactly this, is that over time, because Microsoft no longer likes SMB v1, but that's the only version that DOS understands, and I was trying to put, for the sake of my work with SpinRite, I wanted to get an MS-DOS or FreeDOS system, either in a VM or an actual physical machine, on my network so that I could easily share files among them.

Well, since DOS only knows about SMB v1, it was necessary to get Windows to speak that. Windows 7, no problem at all. I was wanting to run SMB v1. So I do what, like when I'm faced with a mystery, what we all do now is you figure, okay, if I'm having this problem, lots of other have had it before me, so you google. And what I discovered was that apparently everybody has a different version of Windows, of Windows 10. So it's no longer the case that, like, there is a version, and the Internet's knowledge base can be used to solve it.

And so what I realized is what Windows 10 has become is sort of a smear over time. It's just, you know, it's a moving target. And so it's a smear. So anyway, I ended up solving the problem, I think, as a consequence of noting the problem in the GRC newsgroup, and a couple people suggested some registry tweaks which may help. Anyway, I ended up actually coming up with a development solution for SpinRite that no longer required me to do file sharing, which is probably just as well because, boy, I mean, sooner or later Microsoft is just going to lower the boom on SMB v1, which they argue is not secure any longer.

But I just wanted to note that it is - it's unfortunate that they won't leave it alone. Now, I mean, for example, when I do the control panel, you used to be able to flip it into a mode where you would see all the little icons, rather than just sort of the summary of things. In older versions of Windows, it's there. I'm on the latest one, Windows 10, 1909, the Enterprise build. And it's gone.

And so, and I'm trying to do it, and it's like, you know? So I think, okay, how do I show all the little icons? Because I'd rather have the granular view that I used to have, and I'm sure that Windows used to have. Windows 10 used to have. And I go online and say, you know, like ask Google a question, and up come lots of people who have different versions of Windows 10 than I do, and they have it. But not the 1909 Enterprise version. Microsoft in their infinite wisdom said, no, we're trying to phase this out, so we'll just take it off of the UI. It's like, I just - okay. Fine.

Leo: I don't know how much I can tell you about this event yet. Some of our - not you, Steve, because you did the LastPass event. We're not going to make you do everything all the time.

Steve: Going to rotate, that's good.

Leo: We're going to rotate. Some of our other TWiT hosts are coming out to St. Louis. It's going to be a lot of fun. And Steve, LastPass does want us to do more events, and I hope we can lure you out of your Fortress of Solitude to get somewhere else out in the country because people [crosstalk].

Steve: I will begin to have deliverables for SpinRite stuff before long, so that'll buy me a little latitude.

Leo: You must have been working your butt off.

Steve: I'm working. Well, yeah. So there's SpinTest is the tool that I was developing where I was able to get some early benchmarks. That's running again. I've got my development environment established. I've been building stuff. So I'm going to first produce a sort of a little more official benchmark. I'm just curious to know whether I'm right that in the seven years since we last did this, the throughput of drives has increased as a consequence of the density going up. So we know that density's gone crazy. But so should the throughput because drives are still spinning at the same speed. So what I'm going to do is I'm going to turn the work I've got into a benchmark so that anyone who's interested can run it on their drives and get a sense for how quickly SpinRite will, well, first of all, how fast their drives go.

Leo: Right.

Steve: And then how quickly SpinRite will be able to run across them.

Leo: I just bought four of the new...

Steve: Crazy drives, yeah.

Leo: 16TB helium drives from Seagate. 16TB.

Steve: And as I remember, the reason they're using helium is that it allows the head to fly closer to the surface.

Leo: Yeah, yeah.

Steve: Because it's a lighter gas than just - than air. So it's a hermetically sealed environment, and wow.

Leo: I can only imagine how much data loss is happening every second on these things; right? The ECC must be just going crazy, or whatever they call it on drives, yeah.

Steve: Yeah, no, it is ECC, and it is no longer something you only use when a defect is involved. It's just part of the process because it's like, I think that was a one. Well, we'll let that get solved when we get to the end and see how everything added up right.

Leo: Wow.

Steve: Okay. Now, this is just looney tunes. But I thought it was fun and worth mentioning, just in passing. Since we've been talking about Windows 10, I thought it was worth mentioning that Richard Stallman's founded in 1985 Free Software Foundation has asked Microsoft, and this is a term I was not familiar with, Leo, to upcycle Windows 7 by releasing it to the public.

Leo: To hoots and howls of laughter from Redmond.

Steve: Exactly. Uh-huh. Like that's ever going to happen. So they created a page and a petition. FSF, as in Free Software Foundation, fsf.org.

Leo: I'll sign it.

Steve: Yeah, in fact, go check how many are there because they were looking for 7,777.

Leo: Oh, they've got more than that. Oh, yeah.

Steve: Last night they were nearly - they weren't quite double that, but they were 13,000 and some.

Leo: 13,365 right now.

Steve: So, okay.

Leo: They want to open source Windows 7, is what they're saying.

Steve: Yes. They want to open source Windows 7. Despite the fact that XP hasn't been, nor has, I mean, the last thing with the - I don't think even the 16-bit OSes were. MS-DOS finally was because...

Leo: And actually that benefited you, though; right? Because you used FreeDOS in your SpinRite 6 distro.

Steve: Yeah. Also though there were, like, leaks of MS-DOS source...

Leo: Yeah, but you can't use that legally.

Steve: No, no.

Leo: The problem is, and we talked about this on Wednesday with Paul and Mary Jo, the real problem, the thing that stops Microsoft, I'm sure they would love to do this, is there's a lot of proprietary code in there. It's not just their code.

Steve: Leo, Windows 10 is Windows 7.

Leo: And that's the other problem.

Steve: I mean, that's...

Leo: You'd be giving away the keys to the current version, yeah.

Steve: Yes. It's not any different. I mean, look at it. You click about three layers down.

Leo: There's a little overlap.

Steve: And you see the same dialog. Nothing changed. It just, yes, they're trying - it's not like they wrote a whole new operating system. It's the same.

Leo: Yeah.

Steve: So, yeah. Anyway, I got a big kick out of this because they said: "Current signers," that is, of their petition. "We've reached the goal, but there's still time to show your support." I think it closes, oh, it closes tomorrow. They said: "Sign the petition before February 5th to stay updated on the campaign." Like this campaign is going anywhere. "On January 14th..."

Leo: Here's the latest. Microsoft still says no.

Steve: Yeah. Here's the latest, Microsoft seems to be ignoring us.

Leo: Yeah, they're not even saying anything. What?

Steve: So they're telling us like we didn't know: "Windows 7 reached its official end of life, bringing an end to its updates as well as its 10 years" - and get this - "of poisoning education, invading privacy" - this is the way to win friends, right, and convince Microsoft that they should release it - "invading privacy and threatening user security." They said: "The end of Windows 7's lifecycle gives Microsoft the perfect opportunity to undo past wrongs, and to upcycle..."

Leo: Oh, lord.

Steve: Which, again, is a term I've never encountered, "...upcycle it instead."

Leo: Well, they don't want to say "recycle Windows 7."

Steve: "We" - this is Stallman and company - "call on them to release it as free software and give it to the community." It's like, what?

Leo: Stallman's gone, by the way. He was ousted. But Stallman's spirit lives on, clearly.

Steve: Oh, my god.

Leo: We call on you, Microsoft.

Steve: It gets better, "...give it to the community to study and improve."

Leo: Honestly, they could; right?

Steve: Yes. Well, the problem is, you know, the bad guys could study it also and find things that are obscure that could be leveraged. Because, again, it's Windows 10. "As there is already a precedent for releasing some core Windows utilities as free software" - and by the way, I think that was Calculator - "Microsoft has nothing to lose by liberating" - it's liberating - "a version of their operating system that they themselves say has 'reached its end.'"

So here's the best part. So in this it says, I'm not making this up: "To the executives at Microsoft." We have three bullet points. First one. I can't even say this. "We demand that Windows 7 be released as free software." As you said, Leo, they fell out of their chairs there in Redmond. "We demand." That's why I always want to make sure that you're centered over your ball for something like this.

"We demand that Windows 7 be released as free software. Its life doesn't have to end. Give it to the community to study, modify, and share." Uh-huh. "We urge you to respect the freedom and privacy of your users, not simply strong-arm them into the newest Windows version." So they must have their own weed that they're, like, smoking. "We want more proof that you really respect users and user freedom, and aren't just using those concepts as marketing when convenient."

And then they finish, now speaking to those who would sign: "We need your help to send Microsoft a strong message." Well, okay, first of all, a strong message would number in the billions of signatures. I mean, again, not that that would be any more effective, not we didn't quite get double the 7,777 supporters that we were looking for. "We need your help to send Microsoft a strong message. We want 7,777 supporters to take a stand with us for freedom."

Leo: Oh, I agree. I don't disagree with them. It would be wonderful. It's just not going to happen.

Steve: Oh, Leo. I mean, it's not even, I mean, it's not even like maybe it would happen.

Leo: They've put out some of it. I mean, they did put out, what, they put the Calculator out in open source.

Steve: Yeah.

Leo: They could do Minesweeper next. Solidarity

Steve: Valuable, valuable piece of intellectual property, yeah. Oh, my. Please put out Candy Crush Soda Saga.

Leo: Oh, I wish they would put that out of its misery, yes.

Steve: Oh, god. Anyway, so they didn't quite double that number. So anyway, just it was so bizarre, so far out in left field that I just want to share that with our listeners. But

it does give me a segue into an interesting piece of, well, a bemusing piece of news about Windows 7 updates. It turns out that the "was supposed to be the last" January Patch Tuesday update for Windows 7, broke something.

Leo: Yeah, something stupid.

Steve: Oh, my god, Leo. I know. Something that apparently...

Leo: We don't have to test this. It's the last one. Let's just put it out.

Steve: But don't you wonder, like, what are they doing that broke it? Okay.

Leo: Yeah. It's all, well, it's spaghetti code, that's why.

Steve: It is. It's a disaster. So something that Microsoft cannot let lie, believe it or not, it broke the desktop wallpaper stretch functionality, which results in a black-as-night desktop when an image is being stretched to fit the desktop. And so the initial announcement of this said that the Windows 7 ESU, the Extended Service Updates people who are paying now, would have this fixed. Turns out apparently that didn't go down very well somewhere. I don't know, somewhere within Microsoft.

Now, in their announcement of this January 14th, 2020 update, Microsoft acknowledged the error under, and they called it "Known issues in this update." And they said: "After installing KB4534310, your desktop wallpaper might display as black when set to stretch." And then, under workarounds they said: "To mitigate the issue, you can do one of the following: Set your custom image to an option other than Stretch, such as Fill, or Fit, or Tile, or Center. Or choose a custom wallpaper that matches the resolution of your desktop."

Now, of course you can also, if you just happen to love the wallpaper you have, you could use a third-party, I mean, like any image-stretching tool, to stretch it to size, and then just use the, what, Fill or Fit or Center, anything but Stretch.

Leo: Yeah.

Steve: So then they said: "We are working on a resolution and will provide an update in an upcoming release" - and here's the change - "which will be released to all customers running Windows 7 and Windows Server 2008 R2 SP1," because you really do want your Windows Server to have wallpaper that is working properly.

Leo: This shows what Windows users think is a showstopper.

Steve: Oh, my god.

Leo: I can't make my wallpaper stretch.

Steve: You broke my wallpaper with the last update ever. So what are you going to do about it? So anyway, as I said, they initially said ESU only. Now they're apparently saying "all." So if they're going to be doing a rollout, what would that be? Tomorrow, or, no, next Tuesday, the 11th, will be the second Tuesday of the month. Maybe they'll sneak in the month's security fixes for free because why not? If we're going to have to - actually, that probably would not require a reboot. One hopes that fixing the black desktop wallpaper stretch function would not require you to reboot your machine. So who knows.

In more serious and interesting news, we have the remote code execution exploit for Windows RDP Gateway has now been demoed. We recently discussed the discovery of what was at the time only a denial of service attack on Microsoft's remote desktop protocol gateway service. And once again we see that when a sufficiently skilled hacker carefully examines nearly any sort of software flaw, it's often possible for that researcher to discover some way to manipulate the vulnerable machine into executing code of their choosing, whether it's code they provide in a buffer, or whether it's code that's already there, and they just cause the machine to jump around and execute little bits at the tails of various existing functions in so-called "return oriented programming."

InfoGuard AG's penetration tester, Luca Marcelli, demonstrated a working remote code execution exploit. The exploit targets what we were talking about recently, this Remote Desktop Gateway, on devices running Windows Server versions from 2012 through 2019, so all of them except the one before 2012, which was the one based on Windows 7, so that's 2008 R2. And we now have a name. Whereas the previous Remote Desktop Protocol, the RDP protocol, the sort of the raw protocol, that was the now-famous BlueKeep. This one is BlueGate, of course, because it is the gateway service. Marcelli said that a blog post detailing how to achieve remote code execution with BlueGate will be forthcoming in the next few days, but that he wanted to be responsible and "wait a bit until people had enough time to patch before releasing this to the public."

Okay. So that's probably reasonable because it was fixed in January's Patch Tuesday, which is only a few weeks hence. Or, wait, no, in the other direction. A few weeks ago. So, you know, we can hope that these systems are going to get fixed. The problem is that at last count there are more than, I'm sorry, no, not quite 20,000 currently unpatched and vulnerable Remote Desktop Gateways accepting connections worldwide with 6,816 of them in the U.S. alone. And since the use of this vulnerable Remote Desktop Gateway is only higher end server platforms, it's very likely that those are valuable networks. So they're from 2012 on.

And a small company is not going to need the gateway, since remember that that's a frontend behind which Remote Desktop Protocol servers stand. So it's probably larger installations, larger corporations, or government facilities. And if they are not applying patches in a timely fashion, and at this point several weeks downstream from it being fixed, they're still vulnerable. There's still a significant count of vulnerable systems. Now and very shortly we're going to have a proof-of-concept disclosure for its exploitation. So I have a feeling it won't be long before we actually see these exploits moved, weaponized, and occurring.

This is a happy story, finally. Google is really making some payouts for their bounties. Compared to 2018, Google doubled the reward payouts for bug bounties in 2019. I have a graph in the show notes that really demonstrates what's been going on from 2015, '16, '17, '18, and '19. And props to Google. I think one of the things we're seeing is that, and as we know, we've been discussing bounties as a viable career for sufficiently talented hackers. Or part-time, see if you can find one in the evening and sharpen your skills in the process. Google gave out $3.4 million in bounties during 2018, and 6.5 million last year.

So last year the reason for this escalation is they launched their Developer Data Protection Reward Program, which was aimed at uncovering data abuse issues in Android apps, OAuth projects, and Chrome extensions. They're looking for any apps that violate Google Play, the Google API, and Google Chrome Web Store extension privacy policies. And privacy violators are not going to win a huge reward. On the other hand, they're easy to find. Depending on the impact of the bug found, researchers may be rewarded up to $50,000 per report.

Also last year, Google tripled its top reward payouts for security flaws in Chrome from 5,000 to $15,000 and doubled the maximum reward amount for high-quality reports from 15,000 to 30,000. The Android security rewards program added additional exploit categories and raised the top prize to $1 million for a full-chain remote code execution exploit with persistence that compromises the Titan M secure element on Pixel devices.

And recall when we talked about this, like we talked about this happening last May, Google had at that time recalled the Bluetooth versions of the chip after discovering a vulnerability that would allow attackers within Bluetooth range to take control of the device. So there's also the Google Play Security Reward Program, which paid out $650,000 total in rewards during this - just the second half of 2019, after its scope was expanded to any app, including third-party apps that had more than 100 million installs. So it has to be a significant app. But even if it's not theirs, they will pay.

And also recall that there is a 50% bonus bump, which we previously discussed. Their security team added in their posting about this last Tuesday, they said, if you achieve the top reward Titan Pixel M exploit on specific developer preview versions of Android, they will add, they said, a 50% bonus. So that raises the top prize to $1.5 million. The discovery of major problems in pre-release Android, of course, is worth much more to Google because they would really love not to have that exploit get out into the field. And they're willing to pay for it. Of course Google's got plenty of cash. So yeah.

And this is its 10th bug bounty anniversary. They began offering bounties back in 2010. During this past 10 years they've paid a total, a grand total of $21 million in rewards to date. And so it's significant that 6.5 million of that 21, so nearly a third, was just in this most recent year. As that graph shows that you've put up a couple times, Leo, it's really accelerated recently.

Leo: Somebody in the chatroom said, well, this just means Google's got buggy software. Let's flatten that, shall we?

Steve: No. Yes. Well, it means that people are looking. I mean, the lesson we've learned is all sufficiently complex software is buggy. I mean, it is so difficult to produce absolutely bug-free software that it isn't cost effective. The space shuttle computer software had the benefit of being simple, but it was incredibly complex because it had to be accurate. So an incredible amount of money was put into making it bug-free. I believe, and we've talked about this before, that eventually we'll get to, and we have been talking about the changes we're beginning to see, where we have so much excess processing power now that we no longer need to be coding in C, where a string is a pointer to a region of memory. And, I mean, that's all it is. You are given a pointer to a region of memory in C. That's about as dangerous as anything could be.

And so as a consequence we see disasters all over the place because somebody figures out how to cast that as a signed pointer, and now you can put a negative value in it and poke around in memory below the string that you are allocated, and stuff like that. So we're going to be moving to languages which enforce safety where C just absolutely doesn't even try. Now, programmers like having that power. Unfortunately, the lesson

we're learning is programmers can't be trusted with that power. Not that they mean ill, but they just make mistakes. And so in the future languages will catch the mistakes, or the language simply won't give you a pointer that you can do anything with. You'll get a handle to a string, and the language will make sure that there's, like, no way possible for you to abuse that handle.

So anyway, I think what we're seeing is we're seeing that Google is incentivizing people to inspect their code and other people's code and are being willing to pay because we've seen also that everyone talks about, oh, the solution is to open source it, then everybody can look at it. Except, yeah, everybody's busy. And so open source code sits there, open but uninspected. And as we've seen, you need to raise money in order to fund an audit of an open source program in order to get people to look at it and go, "Oh, look what I found over here. This is really bad." And then people go, "Oh, you're right," and then it gets fixed. But the fact that it's open doesn't mean that it's going to get fixed. It's got to be inspected in order for it to get fixed.

And speaking of inspectors, Leo, we have the return of Roskomnadzor. Roskomnadzor is, as we know, Russia's telecommunications watchdog. They announced last Friday that it has instituted administrative proceedings, which sounds kind of ominous, against Facebook and Twitter because of each company's continued refusal to move the data of their Russian users onto servers located inside that country's borders. What could possibly go wrong?

Roskomnadzor said: "These companies did not provide information on meeting the requirements for localizing the databases of Russian users of the corresponding social networks on servers located in the Russian Federation, as provided for in Part 5 of Article 18 of the Law on Personal Data No. 152-03." Bureaucracy much? "Administrative proceedings," they continued, "have therefore been instituted on the grounds of an administrative offense in accordance with Part 8 of Article 13.11 of Administrative Code of the Russian Federation, which provides for an administrative fine in the amount of" - uh-oh, wait for it - "1 million to 6 million rubles." Which is somewhere between $16,000 and $94,000 U.S.

Leo: Oh, it sounded like more. I'm disappointed.

Steve: So Leo, do you think maybe Facebook and Twitter can deal with a fine not to exceed $94,000?

Leo: I guess so. Geez Louise.

Steve: I know.

Leo: Boy, they're really coming down hard on them.

Steve: Whoa, boy, they lowered the sickle on Facebook and Twitter. As we've covered before, Facebook was previously threatened with a ban in September 2017 for the same reason. Twitter agreed to the demands of Russian officials at the time and proceeded to inform the Roskomnadzor that it was planning to move Russian users' data by mid-2018. According to the Moscow Times, Roskomnadzor said Friday that a complaint will also be filed - oh, Leo, they're going to complain - will also be filed in Russian courts next week.

And a new law, just signed last month by Vladimir Putin, imposes higher fines of up to 18 million rubles, Leo, for repeat offenders. That's $280,000.

Leo: Goodness. Terrible.

Steve: Meanwhile, last Wednesday, the ProtonMail Twitter account tweeted that "The Russian government has blocked ProtonMail and ProtonVPN within Russia. We are reaching out to the appropriate authorities to get the block lifted as soon as possible." So this ban was prompted by Proton's refusal to register their VPN services with Russian authorities, which was asked of all VPN providers operating in Russia. ProtonMail and ProtonVPN users are advised by the company to access the two services through Tor.

So, Leo, what's going on? So does Russia really want social media companies operating inside Russia or not? Do they really care where their citizens' data is stored? And if yes to either, then why are years passing with fines so low relative to the sizes of those companies? I just, you know, doesn't make any sense.

Leo: The GDPR fines are percentages of revenue.

Steve: Yes.

Leo: And significant percentages.

Steve: That will get somebody's attention. So this is just nuts. I mean, like what are they actually doing? Are they just wanting to look like they're caring? Certainly they don't want Facebook and Twitter to pull out of Russia.

Leo: They look like they care, yeah.

Steve: Yeah.

Leo: You nailed it, yeah.

Steve: 18 million rubles, ooh.

Leo: Yikes.

Steve: But those rubles are itty-bitty rubles.

Leo: I bet, though, that until they did the currency conversion, there were some people who blanched a little bit at Facebook and Twitter, going, "Oh, that's - oh."

Steve: Yeah. Whoa, they're going to shape up now, baby. And apparently you can just, you know, you get fined, and you ignore it until you get blocked.

Leo: Right.

Steve: And again, are they going to block Facebook and Twitter? Seems unlikely.

Leo: Right, yeah.

Steve: But Facebook did get fined, but not by Russia. Facebook just lost an interesting class action lawsuit it had been fighting for the past five years. And it will pay $550 million in settlement of that suit. So that's, okay, still Facebook. But $550 million, more than half a billion dollars. So that's a significant chunk of money. The lawsuit was brought against Facebook for scanning their users' faces in photos and offering tagging suggestions. The plaintiffs claimed that the platform, Facebook, violated the strictest biometric privacy law in the land, Illinois's Biometric Information Privacy Act (BIPA) as a consequence of its tag suggestions tool.

Facebook started using the tool in 2015 to automatically recognize people's faces in photos and suggest to their friends that they tag them. And it does this without users' permission and without telling them how long it will hang onto their biometrics. The lawsuit contends that Facebook squirreled away face prints in what Facebook has themselves claimed is the largest privately held database of facial recognition data in the world. It's like, yes, let's shout that from the mountaintops. What could possibly go wrong?

Last September Facebook said it was ending tag suggestions in favor of the multipurpose face recognition setting, which it made available to all users, along with an opt-out option. The New York Times in its reporting of this lawsuit outcome, this $550 billion, referred to the 550 - I'm sorry, 550 million hit as "a rounding error" for Facebook. Facebook reported that its revenue rose 25% to 21 billion in the fourth quarter, that is just the fourth quarter of last year, while profit increased by 7% to $7.3 billion. So in the last quarter of last year, it had revenue of $21 billion and profit of one third of that, 7.3 billion.

Yeah, so Facebook also is - they are the new Intel. They are printing money. Although this is a lot of money, that is, the fine, it turns out it could have been worse because Illinois's BIPA requires companies to get written permission before collecting a person's biometrics, whether they are fingerprints, facial scans, or other identifying biological characteristics. I mean, it is broad. It also gives Illinois residents the right to sue companies for up to $5,000 per violation, which could get pretty expensive.

So, not surprisingly, Facebook fought this lawsuit tooth and nail. In 2016 it tried and failed to wriggle its way out by saying that its user agreement stipulates that California law would govern any disputes within the company. On the other hand, I don't know how long they're going to be safe in California. And besides, Facebook said in its motion that BIPA didn't apply to Facebook's facial tagging suggestions for photos. The judge's response was nope on all counts. Going by Illinois law was just fine, the judge said, and it was always clear that BIPA would cover faceprints because it governs the use of all biometrics.

So, you know, we're sort of in a position where we're trying to figure out exactly who owns our biometric data these days. 23andMe just laid off 100 employees, which was

14% of its workforce, as consumer demand for its DNA testing kits has dropped significantly. The company said, in announcing this, that a variety of factors, including privacy concerns, could have contributed to the slowing market. And this of course followed a bunch of news coverage indicating that the company was responding to law enforcement queries of its DNA database.

Since I was curious, I looked up 23andMe's own disclosure titled "How 23andMe responds to law enforcement requests for consumer information." I have the link in the show notes for anyone who's interested. They said in that page, in their own page on their site: "We work very hard to protect your information from unauthorized access by law enforcement. However, under certain circumstances" - now, our listeners know how to hear this. It's very much like Intel saying, "It is possible." Uh-huh.

So 23andMe says: "Under certain circumstances your information may be subject to disclosure pursuant to a judicial or other government subpoena, warrant, or order, or in coordination with regulatory authorities. If such a situation arises, we have to comply with valid governmental requests, and we will notify the affected individual(s) unless the legal request prevents us from doing so." In other words, they're saying we must comply with a lawful court order and its likely accompanying gag order. So, yeah.

And apropos of this, a New York-based facial recognition startup named Clearview AI has amassed a massive database containing more than three billion images scraped from employment sites, news sites, educational sites, and social networks including Facebook, YouTube, Twitter, Instagram, and Venmo. And since this data collection was done without the consent of the people whose images have been collected, that company is now also being sued in what will likely become a class-action lawsuit. Since that suit is citing the BIPA Illinois regulations, the complaint against Clearview AI, this company based in New York, was filed in Illinois.

The New York Times published an expos about how Clearview has been quietly selling access to faceprints and facial recognition software to law enforcement agencies across the U.S., claiming that it can identify a person based on a single photo, revealing their real name and far more. The New York Times said: "The tool could identify activists at a protest or an attractive stranger on the subway, revealing not just their names but where they lived, what they did, and whom they knew."

Clearview told the Times that more than 600 law enforcement agencies have started using Clearview in the last year, and it's sold its technology to a handful of companies for security purposes. Clearview declined to provide a list of its customers. Eric Goldman, the co-director of the High Tech Law Institute at Santa Clara University, told the newspaper that "The weaponization possibilities of such a tool are endless." He said: "Imagine a rogue law enforcement officer who wants to stalk potential romantic partners, or a foreign government using this to dig up secrets about people to blackmail them or throw them in jail."

The New York Times headline about Clearview AI suggested that the company "might end privacy as we know it." From their report they said: "Even if Clearview doesn't make its app publicly available, a copycat company might, now that the taboo is broken. Searching someone by face could become as easy as googling a name. Strangers would be able to listen in on sensitive conversations, take photos of the participants, and know personal secrets. Someone walking down the street could be immediately identifiable, and his or her home address would be only a few clicks away. It would herald the end of public anonymity."

Welcome, Leo, to our sci-fi future. And it is, I mean, it will change our world. I have nothing to hide, and when I'm out in public I'm not in disguise. But since I'm not a high-visibility celebrity, thank goodness, there's a sort of assumption that people who I don't

know, don't know me, just as I don't know them. But it would change if there was an app like "Super Shazam" or "Photo Shazam" where we could point our smartphone camera at anyone and instantly obtain their name, address, full biographical background, and links to their various social media accounts. It would change the world.

Leo: Yeah.

Steve: And I have to say it seems like it's going to happen.

Leo: Oh, I think it happened.

Steve: Well, yes.

Leo: That's the point.

Steve: It did happen for private purchase. I'm sure it costs a pretty penny. But it seems like it's just a matter of time before we end up with an app on our phone that can identify anybody that it sees unless it's regulated out of existence.

Leo: You may remember Google has had that capability for some time and opted not to put it in the phone because they realized what a problem it would be. I mean, I think that technology is out there. The sad thing is, as individuals, we don't have it. But law enforcement has it. Everybody else has it. Businesses probably have it.

Steve: Right, right.

Leo: Our government has it. So we're the only ones who don't have it at this point.

Steve: Yeah, there was some - I skipped over a little bit here. It says the lawsuit claims that Clearview isn't just selling this technology to law enforcement, it's also allegedly sold its database to private entities...

Leo: Of course it has.

Steve: ...including banks and retail loss prevention specialists. So that says that cameras in retail locations are profiling everyone who walks in the door.

Leo: You're in public, dude.

Steve: Yeah.

Leo: I mean, the pharmacist may look at everybody who walks in the door and with his limited abilities try to identify everybody. Delta Air Lines has been testing face recognition identification at gates. And something like 99%, even though they're given an option to opt out, of customers choose it because it's faster. It's a little bit faster. People, you are in public.

Steve: Yeah. And it's true you have no expectation of privacy because, again, you're in public. You're able to be seen. But there has always been some expectation of relative anonymity.

Leo: That's before computers.

Steve: And that's what this changes.

Leo: Yeah. I mean, every bank has had cameras forever, and every grocery store has had cameras forever. The difference is now they're automated, and they know who you are when you lift that package of Tang up off the shelf. I mean, that's the whole purpose of the Amazon Go Store. You don't pay because it just watches you pick stuff up and go out the door. And then they do this pro forma, oh, scan your Amazon - just in case we aren't sure who you are. But they don't need to do that, obviously. I don't know. I mean, I think it's happened. I think that the horse has left the barn.

Steve: Yeah, well, so all of our listeners, get ready for this brave new world.

Leo: Yeah.

Steve: You won't need to, like you'll be able to walk into Starbucks.

Leo: Oh, hi, Leo.

Steve: And they'll - yeah.

Leo: Good to see you. We have your venti latte ready. All right. On we go.

Steve: So we all remember that Avast and AVG were found to be massively tracking their users behind their backs. Their cover story for that tracking was that the users' browsing was being sent back to the mothership so that the URLs could be checked for safety, even though all other browser-based safety checkers offer exactly the same service locally and without sending the user's URL clickstream off to some central database somewhere.

What happened was that the tracking turned out to be quite extensive, essentially sending back all of the user's actions, such as changing tabs, scrolling the page to provide detailed, very detailed monitoring of their users. The data was supposedly anonymized, but the researcher who discovered it, who was the well-known author of

Adblock Plus, felt that deanonymizing the data wouldn't be difficult, and that turned out to have been shown later.

But that wasn't the point or the concern. What was creepy was Avast had also purchased a data analytics firm that was apparently found to be reselling this data to whomever wanted to troll through it. Mozilla reacted to the news by immediately ousting Avast and AVG's browser extensions from their browser repository, their add-on repository.

We subsequently heard that Avast and AVG had dramatically reduced the amount of data being collected. And I wondered aloud on this podcast how their previous business model could be sustained if so much less data was being collected. Well, the other shoe has dropped. It turns out the answer is it cannot be. Avast is winding down, they said, their subsidiary named Jumpshot, following this investigation into the sale of their users' data to third parties that may pose a risk to their users' privacy. Last Thursday Avast said they will no longer have access to user information harvested from users of Avast products and services, which will be fully terminated. And I was curious. Go to Jumpshot.com, Leo, www.jumpshot.com. And we will see what winding down an organization looks like.

Leo: Oh, my god. That's - geez, Louise. Okay. At least they still left the server up. I didn't get a 404.

Steve: Uh-huh.

Leo: Wow.

Steve: Leo is just seeing what I saw last night. Jumpshot has ceased operations, period. Thank you.

Leo: I hope they're sincere about this. They sure got a lot of heat for it.

Steve: Well, they did. A joint investigation conducted by Motherboard and PCMag, which was just published, revealed that information scraped by Avast from their users was handed over to Jumpshot and linked to individuals through a unique ID in an effort to anonymize, yet track them. It turns out it's possible to pick apart data strings and deanonymize users to reveal their identify, tracing their online footprint, browsing habits, and their purchases. In a blog post, Avast's CEO said that the recent news about Jumpshot "has hurt the feelings of many of you" - he's writing to us collectively.

Leo: Hurt, you know, hurt my feelings. Really?

Steve: It hurt, yeah.

Leo: Little more than that.

Steve: Uh-huh.

Leo: It's not about feelings here.

Steve: "And raised a number of questions." As the chief executive, he feels "personally responsible." Well, he is.

Leo: Uh, yeah.

Steve: For the turmoil. Yeah. And apparently, reportedly, Jumpshot was bragging they had access to information from over 100 million devices. And I guess Avast and AVG were very popular free AV offerings. And as we said, is there any such thing as a free lunch? So, but speaking of third-party antivirus, to no one's surprise, every single antivirus vendor has formally announced that they will continue offering and fully supporting their existing AV solutions for Windows 7 for at least the next two years through 2022.

Now, I saw a tweet which suggested that that was only for extended service users, but I don't understand - or ESU, Extended Security Updates. I doubt that that's the case. So ZDNet had a very clear reporting that stated that they're going to be continuing to offer Windows 7 AV support. And it represents a huge marketing opportunity because I, for example, I'm not sure what I'm going to do. I love my Windows 7 setup that I've got. It's mature. It's not that old because it wasn't - it was only last year that XP died on me, and I thought, okay, well, fine. I'll go to 7. I'm not going to 10. I have a Windows 10 setup at my other location which I'm using, and I like it also. But that's the one where I realized that there wasn't a Windows 10 any longer, there were 10 of them. So, yeah.

Anyway, for what it's worth, all the vendors will be continuing to offer AV. And I wanted to ask you, Leo, do we have a recommendation for like if you can't use Windows Defender anymore?

Leo: No. The recommendation is don't use AV.

Steve: Right.

Leo: Right? I just - I find it hard to find a compelling reason why anybody should be - on Windows you've got AV. Defender is as good as anything out there.

Steve: Yes. And I have to say the only time mine has ever made any noise is when it has found a folder of known viruses.

Leo: Well, that's good.

Steve: Yeah. And it's like, oh, okay, good.

Leo: It works.

Steve: Like it's awake, yeah.

Leo: I think we focus on behavior, keeping your system up to date, your behaviors. And both macOS and Windows have pretty good security features in place to keep the most egregious stuff out.

Steve: Well, and our browsers are going to stay supported, and our browsers are doing a very good job. Our browsers are the way this stuff gets into our system these days.

Leo: Right. Well, that and email. And in fact Gmail is really good for that. They filter out almost everything. We use a different - we use both Gmail and another system. So I shouldn't say we, I mean, we use lots of perimeter protection. But none of the machines we use have antiviruses on them because, well, we know. There's a lot of reasons. Not just privacy. They open up holes. They operate at ring 0.

Steve: Oh, my god, yes.

Leo: So they open up holes.

Steve: They cause Windows updates to fail when Microsoft changes a hook that the AV had to hook into. It is a privacy concern because it needs to inspect into your encrypted traffic. And so that's not a supported function.

Leo: I also tell people the other problem with antivirus is no antivirus is perfect, of course. In fact, on average, at best 50 to 60%. And it doesn't catch anything new, which is a lot. So it gives you a false sense of confidence. So I think people, you know, there are a lot of people out there, I've got an antivirus, I'm fine. I don't have to do anything else. And that's really bad. You want to be kind of cautious.

Steve: That's the best thing.

Leo: Yeah. And people always disagree with me. There's a guy in the chatroom now. No, do not use - I say this on the radio to everybody. Do not install an antivirus. And really don't - you don't need one on a mobile device, really. That's even more the case.

Steve: Right, because they're so well...

Leo: They're so locked down.

Steve: They are so well locked down.

Leo: Do you disagree?

Steve: No, I don't. I mean, my Windows Defender will now have stopped being updated, and I don't care. And whatever I had on XP, the previous iteration, it stopped being updated. I don't care. I just, again, I'm just not doing anything. Again, it is a function of who you are. If you're living in the dark underbelly of the web, then I would do my, I mean, then I wouldn't trust an AV anyway.

Leo: Well, that's the point.

Steve: Because of all the reasons you just said.

Leo: Yes, yes.

Steve: There I would do my work in a virtual machine.

Leo: A sandbox, yeah, yeah.

Steve: Yes, in a really strong sandbox, and be very careful with the stuff I download.

Leo: I'm completely comfortable saying on the radio, "Do not use an antivirus." I say it again and again. It's the wrong thing to do.

Steve: I think you're right.

Leo: They're crap, frankly.

Steve: Yeah, they are. An update on the - well, and they're intrusive. I mean, how many times have we been talking about the problems created by AV recently, hooking into the kernel and causing problems.

An update on the WireGuard VPN in the Linux kernel. The lean-coded, fast, modern, and secure WireGuard VPN protocol has made it into the Linux kernel, as Linus Torvalds merged it into his source tree for v5.6. That's the next Linux kernel expected to be formally released in a couple of months. And the list of kernel changes...

Leo: I think this is fascinating. I can't wait to see this.

Steve: Yeah. In the list of kernel changes, the WireGuard VPN was the first thing on the list. So just to remind our listeners, the WireGuard protocol and its implementation is a project from security researcher and kernel developer Jason Donenfeld. He created it as an alternative to IPSec and OpenVPN. In its current form, WireGuard has about 4,000 lines of code, versus the - get ready - 100,000 lines of code that OpenVPN has been dragging forward. And WireGuard doesn't require OpenSSL, which is another blob of huge ancient code. Four thousand lines of code versus 100,000 lines of code.

Compared to current OpenVPN's kitchen sink, "something for everyone" options, WireGuard relies upon a small set of carefully chosen modern cryptographic primitives that are stronger, perform better, and have been highly scrutinized recently by the cryptographic community. It uses ChaCha20 for symmetric encryption, authenticated with Poly1305, using RFC 7539's AEAD construction - that's the [Authenticated Encryption with Associated Data]; Curve25519, which is what SQRL uses, for elliptic curve Diffie-Hellman key agreement; BLAKE2 for hashing and keyed hashing; SipHash24 for hash table keys; and HKDF for key derivation. So all latest state of the art.

WireGuard provides perfect forward secrecy, protection against denial of service, brute force attack protection, key impersonation protection, replay attack protection, as well as support for an additional layer of symmetric key cryptography to offer post-quantum crypto resistance. It's in there now. This lean selection of crypto primitives deliberately dropped all of the unnecessary choices for encryption, key encryption, and hashing algorithms. And that increases interoperability since there's less to mismatch, while also increasing the risks from the support of obsolete crypto. That is, OpenVPN and OpenSSL are dragging all of this old stuff forward because, well, maybe you'll want to hook up to a server that offers that.

The good news is WireGuard never had it, which is the benefit of starting over from scratch. It's also faster because it lives in the kernel space, meaning they're not having to do ring transitions constantly. And because it's so small, it's much easier to audit for security vulnerabilities. It's simple to configure and deploy. There is a 20-page whitepaper on the protocol. I've got a link to it in the show notes. I have not yet made time to study it. But as I mentioned at the beginning of the year, when we talked about this, it's on my list of things to get to.

And it's not just for Linux. It's also fully cross-platform, with implementations for Windows, macOS, BSD Unix, iOS, and Android. There are some VPN service providers for it. Mullvad, AzireVPN, IVPN, VPN.ac and TorGuard are currently supporting the WireGuard protocol with servers. And allow me to reiterate that while I was in Gothenburg, Sweden for the SQRL OWASP presentation there, I happened to meet and break bread with the founder of Mullvad VPN.

And our listeners will remember how surprised I was when I signed up for Sync.com, and the "Send me an email for password recovery" checkbox was unchecked by default. I thought, wow, I couldn't believe that they were, like, choosing the secure option. Similarly, that's the way I felt over dinner as I was chatting with the guys from Mullvad, one of them being the founder and owner as they described their company's philosophy and shared anecdotes about the unforeseen consequences of deliberately wanting to know nothing about their customers, even their IP addresses, which they never log. Anyway, I have a feeling - I've not made the switch yet myself. I'm using SSH a lot, and OpenVPN hardly anymore. But when it's time next to fire up a VPN, I think I'll be looking at WireGuard.

Leo: So it's both client and server?

Steve: Yes, yes. It's considered a secure encrypted tunnel. So the server is the end of the tunnel that listens, and the client is the end of the tunnel that initiates the connection. But otherwise it's just a tunnel, so you're able to tunnel - it's a Layer 3 tunnel, so it operates over UDP. Very cool.

And we conclude with this week's Best Hack of the New Decade. Oh, my god. So the question we pose, why would an ingenious German performance artist place 99 active cell phones into a little red wagon and walk them around the streets of Berlin?

Everybody, let's think about that. Why would an ingenious German performance artist place 99 cell phones, operating cell phones, into a little red wagon and patiently walk them around the streets of Berlin?

The answer? Because it brilliantly confused and spoofed Google Maps into believing that there must be a massive traffic congestion, since so many "cars" were all moving so slowly through a region of the city. There's a YouTube in the show notes, and I snapped a picture of it from his website, showing him all alone. And it's not clear what time of day it is. But, I mean, the sun's up. There's a big shadow being cast by him, and it's not the moon. He is alone on the street because Google Maps shows an absolute red in both directions for a good distance.

And so maybe everybody thought, oh, shoot, I can't go the way I normally do. Apparently this bridge is, like, gone. So I'm going to have to go somewhere else and drive around. So he's walking these streets, during the day, and there's no other cars. There's no traffic. Because apparently everyone's using Google Maps now, and it's like, oh, shoot, got to find some other way to go. Incredible. Anyway, I just thought that was such an amazing hack. And the video shows his little red wagon with this, like, just filled with 99 cell phones which he's all by himself just walking through the streets of Berlin, and Google Maps is basically not functional.

Leo: They each had their own SIM, so they were unique. And he tried stopping, but he said Google Maps wasn't affected unless I move. You have to have some movement.

Steve: That makes sense, that they would - well, because if you're in a restaurant or stopped on the side of the road, no movement doesn't indicate, like, there's still traffic.

Leo: Right, you're just stopped, yeah.

Steve: Yeah.

Leo: He also said - and, see, I think we're getting a little more detail in the video. He had to bring the cart up and down the same street many times. He says also when another vehicle drove past using Maps, Google's system said, oh, there's no traffic jam, and set the street back to green.

Steve: Oh, nice.

Leo: So Google was kind of actively doing this.

Steve: Doing everything it could to filter out any extraneous reporting.

Leo: 9to5Google talked to a spokesperson from Google and got this statement: "Whether via car or cart or camel, we love seeing creative uses of Google Maps, as it helps us make Maps work better over time. Traffic data and Google Maps is refreshed continuously, thanks to information from a variety of sources, including aggregated, anonymized data from people who have location services turned on." By

the way, he got the idea for this after seeing the Hong Kong protests and seeing Google Maps affected by that, even though probably most of those protesters weren't running Maps at the time.

Google goes on to say: "And contributions from Google Maps community. We've launched" - here we go - "we've launched the ability to distinguish between cars and motorcycles in several countries including India, Indonesia, and Egypt." That's because motorcycles traffic split and cut through and drive around, so they go a lot faster.

Steve: Yeah. And so that was generating false non-traffic reports.

Leo: Yes. They say, they go on to say: "Even though we haven't quite tracked traveling by wagon, we appreciate seeing creative uses of Google Maps like this." So they took it in the spirit in which it was intended, instead of saying how dare you.

Steve: Very cool.

Leo: It's interesting. I mean, you learn a little bit about what Google Maps is doing. And yeah, somebody said, well, how do they know it's not 100 people? I don't know. I don't know. There's a lot of heuristics going on in all of Google's stuff. There's gate and so forth. They're looking at a lot of things.

Steve: And we know that heuristics false positive and false negative. I mean, heuristics are like rules of thumb by nature. So it's doing the best job it can. I mean, and really, when you think about it, it's entirely inferential.

Leo: Of course, yeah.

Steve: It's inferring something about the real world from the collection of data that it's getting.

Leo: A subset of the real world.

Steve: And it's not always right.

Leo: Yeah. Although I think it's very telling that as soon as a car goes by, it goes, oh, that's not a traffic jam. So no doubt as soon as this guy stopped going up and down with his little red wagon, those streets went back to green.

Steve: Blink.

Leo: It's fun, though, isn't it. And it's a great video to watch him pulling his little wagon. Steve, always a pleasure. Thank you so much for doing this show. I look forward to it all week long. I know our listeners do, as well.

Security Now! reaches the air every Tuesday. We actually let you watch live, if you'd like to, just as long as you understand that the time is a little soft. We try to get in here around 1:30 p.m. Pacific on Tuesdays, 4:30 Eastern, 21:30 UTC. But "try" is the operative word here. We're usually a little bit later than that because it's the third show of the day. So, but, you know, watch all the shows. Why not? The live audio and video streams are always available, 24/7, at TWiT.tv/live. Most of the time it's recorded content, but during a live broadcast you'll see the live show. You could also get on-demand versions of the show.

Steve's got his own copies, 16Kb audio. It's the only place you can get that. It sounds like Thomas Edison on his Victrola. "Mary had a little lamb." There's also a 64Kb audio, sounds normal. And there are - actually, it's the only place you can the transcripts which Steve commissions. And I think that's a really nice service. Thank you, Steve, for doing that, because for a lot of people reading along is the best way to participate.

If you're watching live, it's nice because you can be in the chartroom, IRC. You hear me talk about them all the time, irc.twit.tv. But we also have forums; Steve has forums. Oh, I didn't mention where you can get all these things Steve has, at his website, GRC.com. GRC, the Gibson Research Corporation, as we were talking about at the beginning of the show, GRC.com, where he is hard at work on SpinRite 6.1. Lots of other free stuff there, including ShieldsUP! and, I mean, just go and check it out. It's a wonderful site, full of goodness.

Our show is available for download on our site, too, TWiT.tv/sn for Security Now!. There's a YouTube channel. I think it's, I don't know, youtube.com/securitynow or securitynowshow, something like that. Actually, just go to YouTube.com/twit, and on the right side all the shows are listed. Each show has its own YouTube channel. Let's see. What else? We have a forum now at TWiT, as well, twit.community, and a Mastodon instance, which is much like Twitter, at twit.social. You're more than welcome to join us in there. Many of our hosts - I'm in there all the time - participate. Thank you, Steve, and I'll see you next week.

Steve: My pleasure. And for the second Tuesday of February, we will see what Microsoft has in store for Windows 7 users.

Leo: Next week.

Steve: Yeah.

Leo: We will also be coming on a little late next week because of the Samsung announcement. We might catch up by then because Samsung usually only does an hour. But Samsung's going to be talking about the new Samsung S20 phone on Tuesday morning. So we're going to push iOS Today forward, push MacBreak Weekly and Security Now! backward, and it should all work out.

Steve: Cool.

Leo: In case you tune in and we're a little late, that's why. Thank you, Steve. We'll see you next time.

Steve: Bye.