SECURITY NOW!

Transcript of Episode #832

Microsoft's Culpable Negligence

Description: This week we look at another very significant improvement in Firefox's privacy guarantees and the first steps for Facebook into native end-to-end encryption. We look at several well-predicted instances of abuse of Microsoft's PrintNightmare vulnerabilities, and at a clever cryptocurrency mining Botnet that optimizes the commandeered system for its own needs. We note ASUS's terrific move to help their motherboard users make the move to Windows 11, and at the merger of NortonLifeLock and Avast. Then, after touching upon a bit of errata and some closing-the-loop feedback from our terrific podcast followers, we conclude with a sober consideration of Microsoft's handling of vulnerability patching during the past year. And we ask what it means.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-832.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-832-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. He's going to explain why the latest version of Firefox is a must-have for anybody who wants to protect their privacy. We're going to talk about Magniber and more PrintNightmares. Also the merger of Avast and NortonLifeLock. Then Steve is going to rip Microsoft a new one. He is mad. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 832, recorded Tuesday, August 17th, 2021: Microsoft's Culpable Negligence.

It's time for Security Now!, the show where we cover your security, your privacy, how computers work, all that, a little science fiction talk thrown in, with Mr. Steve Gibson. Hello, Steve.

Steve Gibson: Leo, great to be with you again for the launch of Year 17.

Leo: Holy cow. Steve, that's something that really should be remarked upon. That's amazing. That's amazing. 17 years.

Steve: I have less hair than I used to.

Leo: That's incredible.

Steve: And you know, if you had any thoughts that maybe Microsoft would ever sponsor this podcast...

Leo: No, I've given up on that long ago.

Steve: It's going to be further away by the time I'm done than it already was.

Leo: Somebody said I hate Windows the other day. And I don't hate Windows. But I'm also a realist about Windows. And this show has probably taught me more than any other show that I do.

Steve: So what's weird about this one is I didn't start off - I explain this a little bit later. But this wasn't my intention. But I was going to do the standard Patch Tuesday Redux, you know, like here's what happened last Tuesday. But some other news came in Wednesday. And I just, you know, sometimes you just have to do a reality check. So this is Security Now! Episode 832 for August 17 titled "Microsoft's Culpable Negligence."

Leo: Wow.

Steve: And I think I'm - and I'm serious about this. I think I've made a very strong case for something being different this year, or maybe in the last nine months or so, at least. Something has happened. And we're on the inside, those of us who follow this podcast, because we see both sides of what's going on. The vulnerabilities and the patches and the exploits, and I guess I was going to say all sides of. But when you put this all together, and look at what's happening, it tells a story which ended up kind of evolving without me intending it to, out of what I was just going to talk about for what happened in August.

But anyway, we're going to have some more fun first. We have a look at another very significant improvement in Firefox's privacy guarantees. Also the first steps for Facebook into native end-to-end encryption. We look at several well-predicted instances of abuse of Microsoft's PrintNightmare vulnerabilities, and then a very clever crypto currency mining botnet that optimizes the commandeered system for its own needs.

We note ASUS's terrific move to help their motherboard users make the move over to Windows 11, and also the merger of NortonLifeLock and Avast. They are no longer separate entities. Then, after touching briefly upon a bit of errata and some closing-the-loop feedback from our terrific followers, as I said, we're going to conclude with what turned out to be a sober consideration of Microsoft's handling of vulnerability patching during the past year, and we have to ask what it means.

Leo: Wow. Lots to come. And a Picture of the Week that we know exactly what it means.

Steve: This one speaks for itself.

Leo: Yeah. Our ever-popular Picture of the Week.

Steve: Okay, now, again, well, obviously a Picture of the Week is always going to be a problem for our audio listeners.

Leo: But you're good at describing these, Steve.

Steve: Yeah, well, in this case - okay. So what we're looking at here is four standard-looking 19" racks, floor to nearly the ceiling. And prominent, and the reason for the picture, is that someone, some, I think the word maybe is "jamoke," what I'm looking for, has hung some 19" rack gear from one corner, four times. Like each one of these 19" racks...

Leo: It's tied up by a zip tie. These are the switches, I guess, 'cause there's a lot of cables going into them, but they're just hanging there.

Steve: And as it would, being hung from a corner, it's not straight up and down, it's off at an angle. And just, you know, what occurred to me is, okay, so the racks were already full.

Leo: Right.

Steve: They look like they're full. But why not set them on top of the rack?

Leo: Well, this is better for cooling, don't you know.

Steve: And I suppose, if you're in a high earthquake environment, they would just swing, like instead of falling off the top of the rack.

Leo: [Tefman] in IRC says, "It's vertical integration, Steve." Wow.

Steve: Oh, boy. So, you know, it is sometimes - remember in the beginning of IP, when the first engineers created those, what were they called, those nodes, I can't remember now, the - oh, shoot, it just escapes me. But they...

Leo: Oh, yeah, yeah, those little elf nodes, or what were they called? They had a funny name.

Steve: Yeah. They did have a funny name. Anyway, they invented the TCP/IP protocol, and like a packet went from San Francisco to L.A. And it was like, <gasp>. It works. Well, looking at this, it's clear we're just taking this for granted now.

Leo: Yes.

Steve: And frankly, you know, maybe it works. What happens when one of those things breaks and falls? I mean, one of the zip ties comes loose. In fact, over there at the far right one, it looks like there's little black things going up in the air.

Leo: Oh, what a mess.

Steve: It's still being suspended; but, oh, boy. Anyway...

Leo: It's also...

Steve: IMP, I-M-P.

Leo: IMPs, that's it.

Steve: Interface Message Processor.

Leo: That's right, yeah. By the way, there's a timestamp on the photo, 2003. I bet you those switches are still hanging there, 18 years later. This is before we even started Security Now!. This goes way back, way back.

Steve: Wow, yeah. Okay. So I think it was the week before last that Paul Thurrott mentioned during Windows Weekly that he had done some testing of the latest Firefox, and he was bullish about it. But I don't recall the details. And Leo, I confess that I may have missed a more recent change of yours. But I think that the most recent browser change you've made as a consequence of what Paul learned was back over to Firefox.

Leo: Yeah. So the week before he made Vivaldi his Pick of the

Week, which is a Chromium-based browser. And I had tried Vivaldi many times. I really like it. You would like it because it's got tabs on the side. I know you like that.

Steve: Ooh.

Leo: Yeah. You could put tabs wherever you want, but you can have them on the left or on the right. And so that's nice. And it's Chromium, so it's fast. It takes probably a couple hours to set it up initially. It's got page after page of settings. Again, something you would like. So I go through this, spend hours, get it all set up, syncing, it's on all my machines. Then the very, very next week, Paul says, "But the browser I recommend and I'm using is Firefox." I was like, oh. Oh. The reason he likes it is the UI. He thinks it's very clean, and he's right.

Steve: They really cleaned it up with that change a few iterations ago.

Leo: Yeah, it's really nice. And it works. And we want to support - as you've said many times, a browser monoculture, at least an engine monoculture, is a bad thing.

Steve: Yes.

Leo: So we want to support them, and they're suffering right now.

Steve: So I have some good news.

Leo: Oh, good.

Steve: Okay. So it's time to catch up on some of the technological changes that Firefox has been making. Because, and this they announced on last Tuesday, so everybody should have it by now, those of us who are still using Firefox - and Leo, I'm glad to know that you're still among the well-washed masses. Last week's announcement is firmly based upon and is an extension of a very significant earlier architectural move that Mozilla made. So I want to first review that innovation upon which last week's announcement was based. And this was back from Firefox 86. And you can remember that number easily because they 86'd cookie tracking with really what was a breakthrough release for Firefox.

Now, unfortunately, Mozilla chose the name Total Cookie Protection, which leaves me a bit cool, since what's that supposed to mean? I mean, at least it's apparently total, whatever it is, so that's good. It's better than if they named it Somewhat Better Cookie Protection. But that still doesn't tell us what cookie protection actually means.

Leo: 72% cookie protection.

Steve: Your cookies are going to be protected? Is that good?

Leo: Yeah, what does it mean? Yeah.

Steve: I'm not sure. I want to delete them. I don't want them to be protected. So brainstorming a bit, a couple ideas came to mind. So I thought, how about if they called it the Total Tracking Terminator?

Leo: I like that, TTT, I like that.

Steve: Wouldn't you rather have - would you rather have Total Cookie Protection or the Tracking Terminator on your side? So even though Firefox's new Total Tracking Terminator has been weakly named Total Cookie Protection, it really does the job.

Leo: That's good enough. That's good.

Steve: Yes, well, and I'll explain why in a second. It would be wonderful, frankly, to see this added into the Chromium core so that all those other non-Google users of the Chromium browser engine, and most of them are, I mean, these alternative browsers are, like, featuring privacy and protection and keeping you safe; right? That's their hook. So it would be great if they could duplicate Mozilla's perfect solution.

Okay. So just to remind our listeners, from February, how does the Tracking Terminator, unfortunately not named that, work? Okay. The best ideas are clean and simple, easy to explain, easy to implement. This is that. Whether cookies are first-party, received directly from the site being visited, or third-party cookies received from any and all other domains whose assets the first-party site has invoked or caused to be invoked, all Firefox does, starting with release 86, when it is set in its strict mode, is to store every cookie received from any domain while visiting a website inside that website's own private what they call "cookie jar." That's all. That's it. That's all that's required for Firefox to effectively terminate all cookie-based tracking.

Third-party cookies still work. They just don't work when you go somewhere else because the third-party cookie you got when you were visiting Site A is unknown to Site B. You may get it again, a third-party cookie for Site B, but Site B doesn't return the third-party cookie you got from Site A. That's the way third-party cookies have always worked. It's always been a problem. And Firefox just killed it. They terminated it by creating essentially stovepiping for the things that your browser gets when it's visiting a site. And I would argue that this is a mistake that every previous browser has made by treating cookies like a global resource which are inherently all shared in a single massive communal cookie jar. And therefore they're accessible from any website.

And this is what this global cookie jar that allowed the same cookie, which may have been planted by one advertising provider, to be used to track the user from website to website as they moved across the Internet. There's always only been a single browser-wide and thus Internet-wide communal cookie jar. But the instant we have individual per-site cookie jars, all that tracking disappears, and you could say that, as I have, tracking is terminated.

Okay. So I'll just repeat that it would be utterly wonderful to see this added into Chromium so that all those non-Google users of the Chromium browser engine - Brave, Edge, Opera, Silk, Vivaldi, and others - could also terminate tracking at the user's choice. I don't think this could be done with an add-on. It probably needs to be implemented pretty deep in the browser's core. But maybe, hopefully, this notion will catch on.

Okay. With that understanding, what happened last Tuesday? The Mozilla security blog posted an update on the privacy-enhancing changes they're continuing to make in Firefox. And now, sadly, they once again gave this new feature the rather milquetoast-y name Enhanced Cookie Clearing. Really? They desperately need someone to snazz up their naming department over there are Mozilla. Instead of the yawn-inducing Enhanced Cookie Clearing name, I was thinking of something more along the lines of Website Historyectomy.

Leo: Oh, no, no, no. No, no, no.

Steve: With Firefox 91.

Leo: No, please.

Steve: Just press the Website Historyectomy button, and Firefox immediately completely and utterly forgets everything, and I mean everything it ever knew about that - now, Leo, come on. You're never going to forget that name.

Leo: No, I remember it now.

Steve: Yeah, yeah. Uh-huh, yeah.

Leo: So what does it do?

Steve: Okay. So Mozilla explains it this way and provides some additional detail: "We're pleased to announce a new, major privacy enhancement to Firefox's cookie handling that lets you fully erase your browser history for any website. Today's new version of Firefox Strict Mode" - and that is of last Tuesday you can get it now, Firefox 91 - "lets you easily delete all cookies and supercookies that were stored on your computer by a website or by any trackers embedded in it.

"Building on the poorly named Total Cookie Protection, what we know now as the Tracking Terminator, Firefox 91's new approach, apparently not to be called Website Historyectomy, to deleting cookies prevents hidden privacy violations and makes it easy for you to see which websites are storing information on your computer.

"When you decide to tell Firefox to forget about a website, Firefox will automatically throw away all cookies, supercookies, and other data stored in that website's cookie jar," which they're still insisting on calling it. "This Enhanced Cookie Clearing makes it easy to delete all traces of a website in your browser without the possibility of sneaky third-party cookies sticking around.

"Browsing the web leaves data behind in your browser. A site may set cookies to keep you logged in, or store preferences in your browser. There are also less obvious kinds of site data such as caches that improve performance, or offline data, which allows web applications to work without an Internet connection. Firefox itself also stores data safely on your computer about sites you have visited, including your browser history or site-specific settings and permissions." In other words, there's a whole bunch of stuff, right, besides cookies.

I'm going to skip some more of this stuff that we already understand about third-party cookies and how they work and so forth. They wrap up: "Embedded third-party resources complicate data clearing. Before Enhanced Cookie Clearing, Firefox cleared data only for the domain that was specified by the user. That meant that if you were to clear" - and they have some examples here about clearing cookies for a specific domain. They would only be cleared for that domain.

They said: "Total Cookie Protection, built into Firefox, makes sure that Facebook.com can't use cookies to track you across websites. It does this by partitioning data storage into one cookie jar per website, rather than using one big jar for all of Facebook.com's storage. With Enhanced Cookie Clearing" - last week - "if you clear site data, the entire cookie jar is emptied for that site, including any Facebook.com data set while embedded in that site."

Okay. So the way to visualize this is that Firefox 91 reorganizes the data that the browser retains. It was originally organized by the domain that owned and produced the data. Now it's organized and also stored by the site you were visiting when the browser received that data. That's a huge difference. What they've done with 91 is extend this stovepiping isolation model beyond just a site's cookie storage. They're now sequestering all history, all data, and all changes that visiting a website can make or cause to be made into a single "Cookies & Site Data" repository, which Firefox's user can delete with the push of a button.

Once this sort of first-party domain-based containment has been provided and as I said, I hope that this is the future of web browsing architecture because it's the correct architecture the only remaining trackable signals being sent by browsers are its static query headers. We've talked about finger-printing based on query headers because they contain version information about add-ons and things - and those of course are in the process of being deliberately blurred now - and also things that can be extracted by JavaScript, like high-resolution battery charge level, current device illumination, GPS location, gyroscope orientation, available storage space and so on. We've covered all those things over the years.

And they are things that JavaScript will send back to a tracker in order to get additional bits of soft identity in order to try to hold onto you. And as we've discussed previously, the resolution provided to JavaScript for each of those real-world physical attributes has been deliberately reduced to blur and increase the entropy of any returned data. I believe that JavaScript should be entirely blinded to all of that without receiving its users' explicit permission. It is not today. You don't see JavaScript having to ask. All of that was only added by techies because they thought it was cool, but not because it serves any clear purpose.

And think about it. How would we feel if advertisements running JavaScript could listen to our microphone or peer out of our camera whenever they felt like it, without our permission? These other parameters are not any really less of a privacy violation. Yeah, it's way lower bandwidth, and it's less directly identifying. But it's of a similar ilk.

So as regards cross-Internet tracking, the handwriting does seem to be on the wall. Users don't like the idea of being tracked. Many may not really care about it that much. But we've seen what the response was to Apple's requiring apps to explicitly request and receive permission to track their users. Nearly everyone said no. If you do it without me knowing, fine. But if you're going to ask me if I actually want to be tracked, then hell no. Are you nuts? And Google appears to be aware that their days of secretly tracking are numbered, too.

So props to Firefox for this change. Essentially what this means is they've sort of picked up on some of the lessons of their incognito mode. That is, you can, without having to go into incognito mode if you want, you can go to a site, muck around there, and then with a push of a button completely remove all previous knowledge from Firefox of your visit to that site. All cookies, all data, you're out of the browser history, no breadcrumbs, no cookie crumbs, nothing is left behind.

So again, all they had to do, the thing I love about this, is that they just switched the store, the global information store, from being global, where it's accessed by domain, they switched it into individual stores by the domain you're visiting, which may contain data from other domains. Those are third parties. But they're all there. And the cool thing is, when you are somewhere else, your browser has no access to that stovepiped information which belongs, logically belongs to the domain you were visiting when you received it. It shouldn't be available when you're somewhere else. That's tracking. And that's the way it's always been done. Firefox fixed it. So yay.

Leo: That actually makes sense. I mean, that was the whole idea of cookies in the first place, was only the first-party site should be able to access that information.

Steve: Yes, exactly. And it was when third-parties began becoming prevalent and said, hey, you know, somebody realized, I can send a cookie? And it'll pop up, it'll be stored there? And when the same guy goes somewhere else? Hey, that's great. Like, yeah, not for us. And so if this change happens, I mean, it is such a clean solution. The remaining problem, as I said, is the fingerprint-y sort of things. And it's well understood. It's recognized. And you have to wonder why a website needs to know the angle at which I'm holding my phone. That's just, you know, yeah, how is that useful?

Leo: Right, right.

Steve: Okay. So something that is useful, Facebook finally adds end-to-end encryption to their base Messenger product. It was back in March of 2019, more than two years ago, that Facebook's CEO, Mark Zuckerberg, proudly stated, as if they had just discovered it, that "The future of communication will increasingly shift to private, encrypted services where people can be confident that what they" - I'm having not to laugh - "that what they say to each other stays secure, and their messages and content won't stick around forever." So of course many industry observers rolled our eyes at this, since even two years ago at this announcement, which was only to express an intention, Facebook was arriving quite late to the party.

Last week, yes, on Friday the 13th, Facebook's Ruth Kricheli, Director of Product Management for Messenger, posted the news: "Today, we're rolling out the option to make voice and video calls end-to-end encrypted on Messenger, along with updated controls for disappearing messages. People expect their messaging apps to be secure and private" - whoa, who woulda thunk - "and with these new features, we're giving them more control over how private they want their calls and chats to be.

"Option for end-to-end encrypted voice and video calls will be available. Since 2016 we've offered the option to secure your one-on-one text chats with end-to-end encryption. In the past, we've seen a surge in the use of audio and video calling, with more than 150 million video calls a day on Messenger. Now we're introducing calling to this chat mode so you can secure your audio and video calls with this same technology, should you choose."

So I did think it was interesting to note that in explaining what end-to-end encryption meant, Ruth's posting said: "The content of your messages and calls in an end-to-end encrypted communication is protected from the moment it leaves your device to the moment it reaches the receiver's device." I don't know how carefully worded that may have been, but at least it was refreshingly accurate. Those who follow this podcast know that the point of attack simply moves interception to before departure or after arrival. And the huge advantage of grabbing it at either end is that you're potentially obtaining all of a targeted user's communications and without the pesky need to filter it out of everyone else's.

And in case anyone listening to this podcast actually uses Facebook's Messenger system - I'm not judging - you might be interested in knowing that they also announced that they had updated its expiring message feature within end-to-end encrypted chats, you know, something I've always been dubious about because presumably you could take a picture of the screen, if you really care. Ruth's post explained that "People don't always want or need their messages to stick around, and the timer controls let someone decide when their messages expire in the chat. We've updated this setting to provide more options for people in the chat to choose the amount of time before all new messages disappear, from as few as five seconds" - boy, you'd better read quickly - "to as long as 24 hours."

And two final points, for the sake of completeness: They also plan to begin testing end-to-end encryption for group chats, including voice and video calls, for friends and family that already have an existing chat thread or are already connected. They're also going to begin a test for delivery controls working with end-to-end encrypted chats. This is designed to prevent unwanted interactions by deciding who can reach the chats list, who goes to the requests folder, and who cannot send messages to the user at all.

And lastly, they plan to launch a limited test between adults in certain countries to allow Instagram DMs to also have opt-in end-to-end encryption. They said, similar to the way Messenger works today, an existing chat or mutual following relationship must exist first. At which point optional encryption can be enabled. So Facebook begins to slowly and cautiously inch its way forward toward this goal of mostly catching up with what other platforms and third-party solutions have long been providing. Better late than never. Eventually they might even turn it on by default. Oh, my goodness.

Okay. So to no one's surprise, and exactly as predicted, attackers have immediately jumped upon the myriad PrintNightmare nightmares to help them move through compromised enterprise networks after first gaining a foothold. Last Thursday, Cisco's Talos group posted: "Another threat actor is actively exploiting the so-called PrintNightmare vulnerability in Windows' print spooler service to spread laterally across a victim's network as part of a recent ransomware attack." They said: "While previous research found that other threat actors had been exploiting this vulnerability, this appears to be new for the threat actor Vice Society.

"Talos Incident Response's research demonstrates that multiple distinct threat actors view this vulnerability as attractive to use during their attacks and may indicate that this vulnerability will continue to see more widespread adoption and incorporation by various adversaries moving forward. For defenders, it is important to understand the attack lifecycle leading up to the deployment of ransomware. If users have not already, they should download the latest patch for PrintNightmare from Microsoft."

And also, last Wednesday the CrowdStrike security blog - and I should mention that CrowdStrike is a sponsor of the podcast - was titled "Teaching an Old Dog New Tricks: 2017 Magniber Ransomware Uses PrintNightmare Vulnerability to Infect Victims in South Korea." They said: "CrowdStrike recently observed new activity related to a 2017 ransomware family known as Magniber, using the PrintNightmare vulnerability on victims in South Korea. On July 13, CrowdStrike successfully detected and prevented attempts at exploiting the PrintNightmare vulnerability, protecting customers before any encryption takes place. When the PrintNightmare vulnerability was disclosed, CrowdStrike intelligence assessed the vulnerability will likely be used by threat actors as it allowed for possible remote code execution and local privilege escalation. This assessment proved accurate in light of the recent incident."

So I just wanted to note that, sure enough, this was the concern, right, that while this didn't provide in the typical case a remote access, even access inside of a network by all the various vectors that we've been talking about the last few months, with what turns out to be a very poorly designed-looking print subsystem for Windows, it immediately got picked up and is being used and being seen now in the wild, leveraged to amplify attacks by anybody who wants to get in and spread throughout an enterprise's network.

Okay. I really got a kick out of this next piece. You've got to love this. A new cryptomining botnet has begun modifying the operating configuration of the CPUs on Linux servers it gains access to, in order to increase the performance and efficiency of its cryptocurrency mining operation. Okay. Specifically, the mining code is modifying one of its processor's MSRs - that's the Machine Specific Registers, which are the - that's the jargon that Intel uses for the registers that are not part of the normal instruction set, but which can be used to tweak the operation down at the instruction microcode level. In this case, the mining code is disabling the CPU's hardware prefetch. Hardware prefetch is an optimization that is enabled by default, as its name suggests. It allows the processor to predict and to load data in its cache, based on the operations that are likely to be required in the near future.

The security firm Uptycs, spelled U-P-T-Y-C-S, spotted a cryptomining botnet that was breaching Linux servers, then downloading the Linux MSR, the Machine Specific Register driver, installing it into Linux, then using that driver to disable hardware prefetching before installing a version of the XMRig which both legitimate and illegitimate users often choose to mine cryptocurrency. Uptycs believes that the attacker likely got the idea to disable hardware prefetching after reading the XMRig documentation, which states that XMRig can obtain a 15% speed boost, that is, improvement in cryptomining performance, if hardware prefetching is disabled.

As we know, hardware prefetching is a good thing, if it's able to properly anticipate the CPU's future needs by prefetching data that does wind up being needed. But that prefetch data needs to be stored somewhere. So prefetching can be a bad thing if the prefetching logic is misfiring, prefetching data that's never used and in the process evicting data that was going to be reused from the processor's memory cache, thus forcing it to be reloaded. So it can work against you.

Well, it turns out it works against the XMRig's cryptomining algorithm. In this case, after infection with this botnet, not only is the infected machine going to run hotter, need more cooling, consume a lot more power, and have its overall lifetime shortened, but it's going to be slower to service the machine's legitimate needs, not only because the CPU will be busy, you know, mining cryptocurrency, but also because when it begrudgingly releases some spare cycles to do whatever it's supposed to be doing, a useful speed optimization, which probably would have been speeding up that process, will have been turned off as a consequence of the infection. So yes, just another thing to look for.

I wanted to also mention, just sort of in keeping track of the industry, that NortonLifeLock and Avast are merging their users and businesses in a deal valued somewhere between $8.1 and $8.6 billion. So to provide a bit of context and history, prior to this merger with Avast, NortonLifeLock had acquired the German AV maker Avira [A-V-I-R-A] for $360 million in cash last December, so December 2020. And almost exactly a year before that, NortonLifeLock was formed in November of 2019 when Symantec sold off its enterprise business to Broadcom for $10.7 billion, then essentially rebranded themselves as a consumer and small business operation as NortonLifeLock.

So now Avast is disappearing, also being absorbed into NortonLifeLock, which will be growing still larger. Under the terms of the merger, Avast shareholders will receive a combination of cash and newly issued shares in NortonLifeLock and will hold approximately 14% of the merged company's total shares. The combined company will have dual headquarters, one in Prague, that's in the Czech Republic, and the other in Tempe, Arizona. So essentially both of the companies' headquarters get to stay. I don't remember who's who.

Leo: Avast is in Prague.

Steve: Oh, yeah, I know that Avast is in Prague.

Leo: Oh, who gets to be president? Yeah.

Steve: Yeah. One gets to be president, and the other guy is going to be the CEO. So kind of a co-president.

Leo: It's kind of funny because both Avast and Norton have had security problems in the past. I think Avira did, too, if I remember correctly. I might be wrong on that.

Steve: Yeah, I do, I think we might have talked about Avast.

Leo: I know Avast had a big issue. So great. So there you go. Three products we don't recommend, all wrapped up into one. But what's amazing to me is $8 billion for Avast.

Steve: Yes.

Leo: Holy cow.

Steve: And here's the other thing. That seems like a lot of money.

Leo: Yeah.

Steve: They're going to have an estimated user base of half a billion users.

Leo: Wow.

Steve: 500 million users.

Leo: I wonder if that includes all the people who've got it installed as trial ware on their Windows install.

Steve: Yeah, exactly.

Leo: Yeah.

Steve: Again, as you said, Leo, and I'm glad you mentioned it, nothing that we're suggesting anybody do.

As we know, the fact that Windows 11 would require support for v2.0 of the Trusted Platform Module, TPM v2.0, this came as surprising and unwelcome news to many Windows 10 users whose hardware lacks TPM 2.0. And this is particularly annoying since Windows 10 runs just fine on such hardware, and we all know that Windows 11 is just Windows 10 with its pointy corners rounded off. Yet Microsoft has said no TPM 2.0, no Windows 11 for you. So in a very nice bit of news from the motherboard manufacturer ASUS, which will hopefully become an industry-wide trend, ASUS is working to make the upgrade to a Windows 11 easier, and in some cases possible for users of their motherboards by offering updated BIOSes for, get this, 207 different motherboards.

Now, some ASUS BIOSes may already support Windows 11 requirement for TPM 2.0 and might only need to have one of two possible BIOS settings enabled, depending upon which processor the motherboard uses. So that would be either you'd turn on Intel Platform Trust Technology (Intel PTT) or AMD Platform Security Processor. That's what they're called in their respective ASUS BIOSes. Turn whichever one of those you see on, and if you have TPM 2.0 you're good to go. But since the BIOS may be somewhere some users fear to tread, an alternative is to simply download and run ASUS's new BIOS for your particular motherboard, any one of 207. It will load with TPM v2.0 deliberately pre-enabled, and thus ready for the forthcoming upgrade to Windows 11, painlessly. I've got a link in the show notes to ASUS's page, or I'm sure you can just find it at ASUS under "Getting ready for Windows 11."

Leo: There's TPM in software, right, like BIOS TPM; right?

Steve: Well, yes, there is TPM in - Intel has a Secure Enclave technology which essentially allows it, as long as the processor supports Secure Enclave, then that's sort of a - it's no longer a separate chip on the motherboard.

Leo: Right.

Steve: It is built into the processor.

Leo: Got it, got it.

Steve: And that's actually better because, as we know, you can hang a digital logic analyzer on the little 8-pin EEPROM which the TPM Module uses. And people have done this and watched the BitLocker key move across the wire and capture it. So it really is better if it's built-in. Although there have been some security concerns about how tight that Enclave is locked up from Intel.

Leo: Right, right.

Steve: But anyway, I thought it was very cool that they're saying, hey, not a problem if you have an ASUS motherboard. Find your motherboard, download the program, you know, it's a Windows-based update. So you don't have to figure out how to boot yourself or anything. And it will update your BIOS, giving you TPM 2.0 if you don't have it, and making sure that it's turned on. So then whenever it is that this happens, Windows 11 will just run on that hardware. So nice going, ASUS. Congrats on that.

A little kind of a funny piece of errata. L.T., although he doesn't think it's funny, L.T. Southall, and he tweeted, I thought this was interesting, his Twitter handle is @oiliswell, O-I-L-I-S-W-E-L-L. And that sort of suggests that he is authoritative. He said: "Every time I hear you or Leo mispronounce the term SCADA, I expect someone will correct you. That obviously has not happened. The correct pronunciation is 'scayda.' The first A is long." And I'm probably guilty of saying SCADA. I think, I mean, that sounds familiar.

Leo: Well, I've got to point out it's an acronym, so the pronunciation is not defined.

Steve: Right. Anyway, he says: "I do know what I'm talking about, having spent 40 years in that business." And given that his Twitter handle is @oiliswell, yes...

Leo: He's in the oil business.

Steve: I get a feeling that maybe SCADA and you are buddies. So thank you. I appreciate the correction. I will do my best to make the first A...

Leo: I'm going to still say SCADA, sorry. I think you can, just like GIF, you can choose your pronunciation of acronyms.

Steve: Yes.

Leo: It's not...

Steve: Exactly. And there again, I've always been a "jif" person, so...

Leo: Supervisory Control and Data Acquisition.

Steve: That's right. That's right.

Leo: I'll pronounce it that way. How about that?

Steve: So I just wanted to mention something that did come back from Darragh Duffy. He said: "Hi, Steve. I just listened to Episode 831. Another suggested reason why Apple are implementing a portion of CSAM on their device is that there's a suggestion that Apple is going to encrypt iCloud backups at some point in the future." He says: "(Currently they don't.)" He says: "Congrats on the 16 years of great work and security insights."

Okay. So the general consensus has settled into it not being the idea of checking cloud-based image storage for illegal content that so much upsets people. And Leo, I know you guys talked about this at the beginning of MacBreak Weekly. We're all trying to not keep talking about this, so I'll just make this one last point. We made the point here on this podcast of noting last week that currently everyone but Apple is already doing that. And the fact that those who are scanning are discovering a truly disturbing amount of illegal photographic content suggests even more strongly that Apple needs to be doing this, too, so as not to become the de facto safe harbor for such material.

But what's being regarded as Apple's mistake is that their solution is to load the hashes of these illegal photos onto everyone's individual devices rather than just quietly doing it themselves on their servers. No one wants to have anything to do with even pre-digested hashed versions of that crap on their personal devices. On Sunday's TWiT show, Mike Elgan made the point that users are wrong to think that they have any sovereign governance over the content and operation of their various devices, since iOS is only available for their use under license from Apple.

Of course, while this may be literally and legally true, and Mike of course is right, it's the perception that matters, and Apple is otherwise all about amplifying and highlighting the fact that these are intimately personal gadgets. Apple probably figures that this will all blow over in time and will just become part of every system's accepted operation. It's going to be interesting to see how this evolves.

I'm skeptical about the motivation Darragh notes about whether this is being done in this way, different from the way everyone else does it, in preparation for Apple making iCloud even further encrypted. iCloud's lack of full end-to-end encryption, if that's what we want to call it, has useful recovery benefits for users. All of Apple's ecosystem has already accepted this aspect of iCloud, and no one appears to care. And it does create a bit of useful safety valve for law enforcement's needs, allowing Apple to respond at least a bit in some specific circumstances to subpoenas.

And it's really not clear to me how much people actually care about encryption. It seems like a good thing. I think that everyone will take whatever they can get. If it's there, great. But people see features and encryption tending to limit what can be done. Sure, super-protecting photos of their cat and last night's lasagna is likely not high up on people's lists. And if it is, don't upload them to the cloud. Fine. These days, most people understand that once something is "in the cloud," it's never possible to really remove it. So anyway, nice piece of feedback. I don't think we'll be talking about this anymore.

Leo: Yeah. Well, okay. Good luck.

Steve: Hopefully.

Leo: We're trying not to, anyway.

Steve: Yeah. In a really great example of one of my favorite concepts that we've talked about on this podcast many times, what I call the "tyranny of the default," Joe Lyon tweeted. He said: "The tyranny of the default. In Germany, there is about a 12% rate of organ donation. In neighboring Austria, that rate is 99%."

Leo: Wow.

Steve: "In Denmark, the rate is 4%; and in neighboring Sweden it's 89%. The Danes and the Swedes are very much similar in almost all respects, so why the huge discrepancy? The answer is that in Denmark, where the donor rate is 4%, you check the box if you would like to opt INTO the organ donor program. Whereas in Sweden, where the rate is 89%, you check the box if you would like to opt OUT of the organ donor program."

Leo: Interesting. The tyranny of the default, yup.

Steve: Tyranny of the default. When you really don't have to otherwise worry about the tyranny of the default any longer.

Paul Durham, he said: "Hi, Steve. You and your incredible Security Now! have inspired me to do something about the time gap between available and applied software and firmware updates, and how this gap leaves the consumer vulnerable. I am building a platform to allow developers and vendors to broadcast the availability of firmware and software updates to subscribers. You frequently mention this gap as a serious issue, and I hope our platform will significantly close that gap. Your thoughts on whether this will be useful and worthwhile? Thanks."

This came via DM. I'm sure Paul doesn't mind me sharing it. I replied, saying go. I think that would be absolutely great. And I think if there were any one thing I can think of that has a hope of helping, it's that. I have talked, I said, like make sure you're getting the email of your vendors. If somebody retired and takes their email account with them, and that's where they were going, email should go to a vulnerability alert account, rather than to an individual, and then be maybe forwarded or copied to somebody who's currently handling it. I mean, I don't know how. But clearly it's important for the word to get out. So I'm anxious to see what Paul comes up with. I just wanted to let everybody know that such work is in progress, and to share my props to him for this.

And, finally, CatGuy tweeted: "Hi, Steve. I have a fairly simple inquiry for you. Given your past episodes on the dangers of end-of-lifed routers and other peripheral devices, is it safe to use such a router as an access point behind a NAT router? I'm starting to segment my home network, planning on using a smart switch and VLANS. All the IoT was going to live on my old 2.4 GHz D-Link. But it occurs to me to question if there is risk here, too."

And to answer your question, CatGuy, I will tell you that is exactly what I have done. I have another WiFi router set up as an access point. It's an older one. Because almost all IoT, I mean, some newer IoT I know also offers 5G, but they all offer 2.4. And like I've talked about these various electric plugs and thermometers and things, and thermostats, they're all 2.4 because that's sort of the universal, and they're trying to keep their cost to an absolute bare minimum. As I said, I don't know how you can sell me a light switch, I mean, a light plug which is a switch, which has Bluetooth and WiFi and a processor, and is on the Internet doing DHCP to get an IP and connecting to servers in the cloud, for $5. I mean, $5. Really. I don't know how that happens. But it does.

And anyway, so yes, what CatGuy has done is my architecture, it's on the bookshelf up there above me is a wireless router running not as a router, but as an access point, and it is plugged into one of the ports of my little pfSense device on a different entire network. I run internally here, I'm on 10-dot. It's a 192.168 network. So completely different network. Sometimes I need to briefly talk to those devices, so I will establish a connection across them in my pfSense config, do what I need to do, and then take it down so that WiFi is completely separate from the rest of my WiFi and wired network here. So that's the way to do it.

Leo: That's your three dumb routers. Well, it's two dumb routers, yeah.

Steve: Yes, the equivalent.

Leo: Yeah.

Steve: Yeah. Actually the dumbness was the way to use NAT. Instead I've got pfSense, which is a super smart router.

Leo: So good, yeah.

Steve: So it's one super smart router

Leo: One super smart. And then one dumb.

Steve: But what's cool is that it's actually four NICs. So each, instead of being...

Leo: So you do have VLANs.

Steve: Yes.

Leo: Actually hardware LANs, really.

Steve: You actually have hardware LANs. Yes, they have separate NICs. So one NIC is a 10-dot, one is a 192.168, and they're just - there's no way for them to see each other.

Leo: Yeah, yeah. Very nice. Very clever. Okay, Steve. Let's hear it. What did Microsoft - what did they do this time?

Steve: Oh. So as I mentioned at the top of the show, I didn't start off today's podcast with this title or topic in mind. Far from it. This section for today was originally up where it usually is, with the generic security news, under the title of Patch Tuesday Redux. But sometimes it's necessary to step back and perform a bit of a reality check. One piece of news from last week hit me as being so unconscionable that, as I started to explore it and what it actually meant, it became clear that the only way to read the facts was that something has gone very wrong at Microsoft. I have no illusions that this podcast will change Microsoft's behavior. But perhaps it's time for us to think about changing ours.

So what follows is what I started off writing, so it starts off sounding like any other Patch Tuesday update. I said: "Last Tuesday, Microsoft released fixes" - it is a Patch Tuesday update - "Microsoft released fixes for 44 security vulnerabilities, with seven of the vulnerabilities being rated critical and three of those being zero-days. The other 37 were rated as being important. Even though the total of 44 is back to being fewer, 13 of the patches fixed remote code execution vulnerabilities, and eight were information disclosures.

"The affected Microsoft products included .NET Core & Visual Studio, ASP.NET Core & Visual Studio, Azure, Windows Update, Microsoft Print Spooler Components, Windows Media, Windows Defender, Remote Desktop Client, Microsoft Dynamics, Microsoft Edge (the Chromium version), Microsoft Office, Microsoft Office Word, Microsoft Office SharePoint, and others.

"Perhaps the most prominent patch released last Tuesday dealt with the Windows Print Spooler Remote Code Execution vulnerability, which has been a major focus since its disclosure in June. And what makes Microsoft's recent performance all the more embarrassing is that the day following Tuesday's patch batch, last Wednesday, believe it or not, Microsoft acknowledged still another remote execution vulnerability in Windows

Print Spooler which it said it's working to remediate. This Print Spooler remote code execution vulnerability is being tracked as CVE-2021-36958 and carries a CVSS score of a mere 7.3."

In their disclosure of this problem, Microsoft wrote: "A remote code execution vulnerability exists when the Windows Print Spooler service improperly performs privileged file operations. An attacker who successfully exploited this vulnerability could run arbitrary code with system privileges. An attacker could then install programs; view, change, or delete data; or create new accounts with full user rights." So no surprise there, right? That's the standard boilerplate for all the bad things that can happen, and also typically do happen, whenever we allow bad guys to remotely execute their code on our machines.

Then in the show notes I have a tweet from Victor Mata. He tweeted on August 11, which was last Wednesday: "Hey guys, I reported the vulnerability in December '20, but haven't disclosed details at MSRC's request. It looks like they acknowledged it today due to the recent events with print spooler." So now here we are again with another newly disclosed Windows Print Spooler RCE. That's not good. I mean, It's really not good. But what's difficult to understand is that we're also told that Microsoft was first made aware of this problem way back in December of 2020 by Accenture Security's Victor Mata of FusionX. So another remote code execution vulnerability in the Windows Print Spooler, which Microsoft has known about since December? And now it's mid-August. And now they're telling us about it and saying that they're scrambling to fix it.

Will Dormann, CERT Coordination Center's Vulnerability Analyst, almost predictably tweeted in that thread that Victor started. He tweeted: "Sometimes I wonder why I bother writing things up and notifying vendors." Yeah, I'd wonder, too.

Microsoft is nothing if not a savvy software publisher with effectively unlimited financial resources. Microsoft's current cash on hand is $130 billion. $130 billion of cash just lying around right now. They could have afforded to hire a talented coder to fix this one problem without even noticing the expense. Not even a rounding error. So they must have decided and I'm really not kidding about this. They must have decided, in some gold-plated ivory tower somewhere, that bugs in their code don't really matter that much.

We all assume, "Oh my god, a remote code execution exploit. The sky is falling." But Microsoft clearly doesn't think so, or they'd prop up the sky if that was a problem. They certainly have the money to do so. But shhh, don't tell anyone. I don't think they really care anymore.

Think about Microsoft's behavior all this year through the Exchange Server fiasco, which directly hurt and damaged so many of their own customers. Not other people's customers. Their customers. Their software. Does anyone think that they lost a single one of those Microsoft enterprise customers as a result? We know they didn't. Microsoft is the only game in town, and the prior investment in Microsoft's ecosystem is far too great. So what did not fixing those Exchange Server flaws quickly cost them? Nothing.

And notice that the attackers are the ones who are increasingly being blamed. We're not blaming the victims of ransomware attacks. We're not blaming the faulty software which those attackers used to gain their foothold, exfiltrate victims' proprietary data, and encrypt their victims' machines. Now we're blaming the attackers. It's their fault for taking advantage of our flaws and weaknesses. It's their home government's fault for allowing them to do that. The U.S. Government is loudly screaming, "You'd better stop attacking us, or else," while Microsoft sits on another remote code execution flaw in Windows Print Spooler for eight months.

Microsoft's not dumb. They didn't get to be where they are by being dumb. Microsoft has always known what matters. And any unbiased appraisal of their demonstrated behavior this year would have to conclude that they are now only paying lip service to their software vulnerabilities, and only then because the politics of the situation requires them to at least appear to care. They are allowing a great many serious software vulnerabilities, of which they have been previously made aware, to remain unpatched for months, while sitting on $130 billion previously paid to them by those same customers who are being directly hurt by those easily patched vulnerabilities.

By delaying the repair of the Exchange Server vulnerabilities at the start of this year, which they were told of in 2020, but didn't bother to repair until they were being used to attack their own customers by the end of March 2021, they directly enabled those devastating attacks against their own paying customers. And so now we're learning of another case where they've known of a remote code execution vulnerability for eight months. How are we to understand any of this, except as the result of a brutal cost-benefit analysis? They have so much money that they could easily arrange to fix these things if they cared at all.

It's not as if these are difficult problems. When it suddenly becomes an emergency, the problems are fixed and released immediately. And it's not as if these are unknown problems. Researchers are bringing these problems to them to fix, asking them to do that, literally handing them to them. But time and time again, Microsoft doesn't bother. Are they so busy working on Windows 11, getting those rounded corners just right, arguing about whether or not to force the new centered menu upon their Windows 11 users, that they can't spare even one employee to fix a serious problem that's been laid at their feet?

At this point in 2021, I think we really need to stop and ask ourselves an important question. Is this the behavior of a company we should continue to support? Is this the behavior of a company that deserves our trust and loyalty? Really, do they have it?

What we have been seeing this year is culpable negligence on Microsoft's part. There is no possible excuse for their behavior. The only possible explanation is that they just don't care anymore. They have the money, they have the resources, and they're being handed the knowledge to prevent devastating attacks against their own customers who have enriched them. And they're doing nothing about it because they don't have to. They have a monopoly on desktop computing. There's no reasonable alternative to Windows. In the past, they used their monopoly to abuse their competitors. Now they're using it to abuse their own customers.

Leo: Ooph. I can't disagree with you. I'm always, you know, trying to understand it, put myself in the other guy's shoes. And I just for the life of me can't come up with a good reason why they would, for instance, wait eight months to patch something, and wait till it's a zero-day, and then say, oh, I guess we'd better fix it. It just doesn't - I can't think of any reason why you might not want to fix it.

Steve: I know. And I've been talking about this. I've been pointing to this this year. It's become insane. I mean, it is not sane.

Leo: Is it possible there are so many flaws, I mean, that there are literally tens of thousands of known serious critical flaws, that they just can't fix them all? They have to triage it and have to wait till it's a zero-day and then say, oh, well, okay, we can fix that one now? Is that possible?

Steve: No, no, because they fixed 44 this month. It could have been 45. But it wasn't.

Leo: It wasn't.

Steve: And they knew about it for eight months. And Leo, $130 billion. They could have a second whole Microsoft that just fixes bugs.

Leo: Yeah, it's really hard to...

Steve: They could. They could.

Leo: ...to think of why this could possibly be. Except just lack of care. They don't care.

Steve: And it's not like they're being forced to discover them. They've got the whole security community finding them for them, showing them. Here's a problem. Here's a zero-day.

Leo: These are good people. I know many of these people. They're not malicious people. They're just like you and me. There's got to be an explanation beyond that.

Steve: There is no cost to them, Leo. It does not cost them.

Leo: So it's a pure business decision. It's not worth spending the time or energy to fix it because it doesn't cost us any money.

Steve: They don't need to. They lost not a single customer.

Leo: Well, they lost me a long time ago.

Steve: And goodwill? They don't have anyone's goodwill anymore. They've got our balls is what they've got.

Leo: Who needs goodwill when you've got them by the short hairs? Exactly. Well, they don't because I think there's Linux, and Linux is growing very rapidly. I'm not proposing that it's more secure, but at least they try to fix the problems as soon as they crop up. What would you, I mean, Steve, you still are going to use Windows because you have to. They do have you.

Steve: Yeah.

Leo: They've got you.

Steve: Yeah. It is the most functional platform. It's not even, when I'm looking for software, many of the things I want just they're not on the Mac.

Leo: Yeah, no, I'm with you on that.

Steve: I mean, it's just...

Leo: I think you would find what you needed on Linux. I'm thinking. But maybe not. I don't know. Linux, I think, fits your ethos a lot closer.

Steve: Well, but the nation's enterprises aren't going there.

Leo: They're not doing it, and that's where your business lies is people buying your software running on Windows. So you don't have a choice.

Steve: Well, and actually my software will be booting on Intel machines, and I will be providing ISOs. I mean, one of the things happening with, well, actually with 6.1 now, but really with 7 is I've become completely OS agnostic.

Leo: Good, good.

Steve: And that's just the reality of the world. But again, Microsoft's not going anywhere. They just don't have to do a better job, and they're not. But it is causing real damage to their customers.

Leo: Well, you're right to call them out. You're really right to call them out. I mean, that is absolutely the case. Do better. You need to do better, Microsoft. Well, there we go. Security Now!. I think you should change the title to "They've Got Our Balls." But, you know, okay, we'll go with "Microsoft's Culpable Negligence." It's not quite as punchy, but...

Steve Gibson is at GRC.com. That's a good place to go to get a copy of the show. He has 16Kb audio, 64Kb audio, and transcripts so you can read along as you listen. That's always great. He's also got lots of other great stuff, including SpinRite, the world's best mass storage maintenance and recovery utility. Lots of freebies. You can leave him comments there at GRC.com/feedback. Probably the better way to respond to Steve or to pat him on the back is @SGgrc on Twitter, @SGgrc on Twitter. His DMs are open, so you can do that, too.

We do the show and record it and edit it and put it up on our site, as well, TWiT.tv/sn. You can get yourself a copy right there. If you want to watch us do it live, we do it on a Tuesday afternoon, try to do it right - we do it right after MacBreak Weekly, aiming for 1:30. Sometimes it's a little later. Today it was 2:00 p.m. Pacific, 5:00 p.m. Eastern, 21:00 UTC. Tune in live and watch. And while you're watching or listening live you can also chat live at irc.twit.tv. Steve, I guess we'll come back here Tuesday, if Microsoft doesn't have you assassinated between now and then.

Steve: As I said, Leo, I think they're not going to be sponsoring this podcast any time soon.

Leo: I can live with that. I absolutely can live with that. But again, you know, that's one of the reasons I like having the idea of "listener supported" because then our only obligation is to you. Anyway, thank you, Steve. Have a great week.

Steve: Thank you, my friend.

Leo: See you next time.

Steve: Right-o. Bye.