# Security Now! #832 - 08-17-21

# Microsoft's Culpable Negligence

## This week on Security Now!

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' 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.

## Sadly, this one speaks for itself...

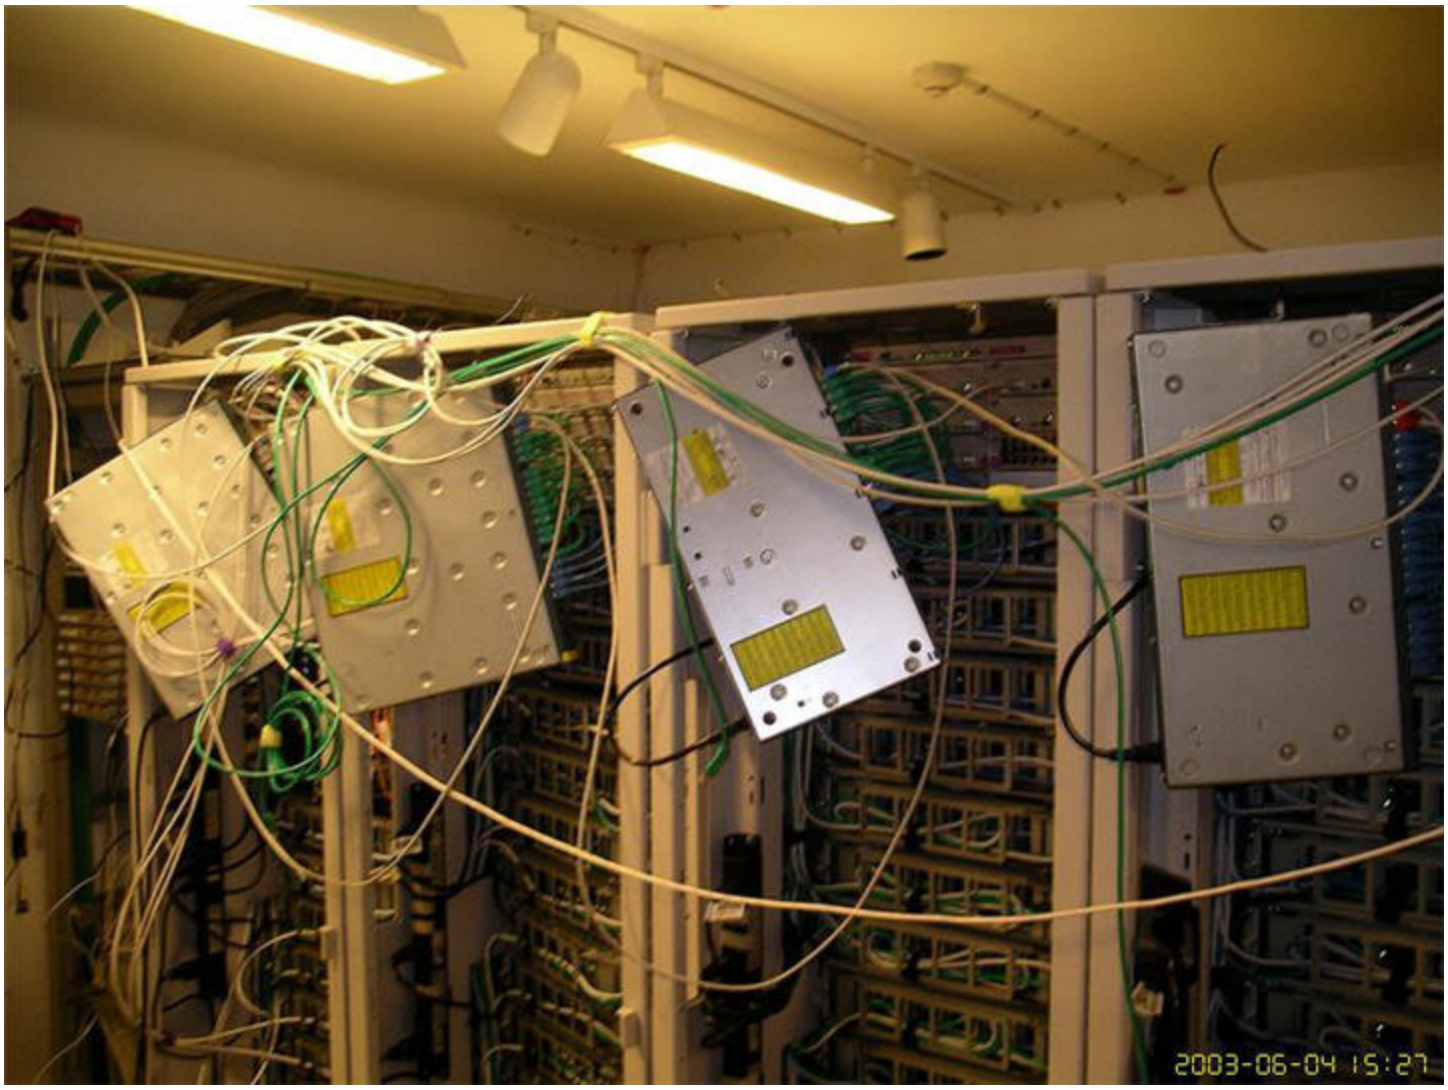


I suppose the racks were full? (They appear to be full.) But why not at least set them on top?

# Browser News

**Firefox Update**

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, but I think that the most recent browser change you've made, as a consequence of what Paul learned, is back over to Firefox, right?  Is Firefox again your default URL handler?

In any event, it's time to catch up on the technological changes that Firefox has been making.

Because last week's announcement is firmly based upon, and is an extension of, a super-significant earlier architectural move that Mozilla made, I want to first review that innovation which was released back in February with Firefox 86. You can remember that number because they 86'd cookie tracking with that breakthrough release of Firefox. Unfortunately, Mozilla chose the name "Total Cookie Protection" which leaves me feeling a bit cool, since... what's that supposed to mean? You know, whatever it is, it's apparently "Total"... 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.  Brainstorming a bit, a couple of ideas came to mind. How about if they'd called it the "Total Tracking Terminator" — Come on. Would you rather have "Total Cookie Protection" or "The Tracking Terminator" on your side? Even though Firefox's new Total Tracking Terminator has been weakly named "Total Cookie Protection," it really does the job. It would be wonderful to see this added into the Chromium core so that all of those other non-Google users of the Chromium browser engine could duplicate Mozilla's perfect solution.

Okay. So, how does Mozilla's "Tracking Terminator" work?

The best ideas are clean and simple. Easy to explain and easy to implement. And this is that. Whether cookies are 1st party cookies received directly from the site being visited, or 3rd party cookies received from any and all other domains whose assets the 1st party site has invoked, or caused to be invoked, all Firefox does, starting with release 86 when it is set in Strict mode, is store EVERY cookie received from any domain, while visiting a website, inside that website's own private cookie jar. That's all. That's it. That's all that's required for Firefox to effectively terminate all cookie-based tracking. The mistake every previous browser has made is treating cookies like a global resource which are all stored in a single massive communal cookie jar. And are therefore accessible from any website. It was exactly this global cookie jar that allowed the same cookie, planted by one advertising provider, to be used to track the user from website to website as they moved across the Internet. There has always been only a single browser-wide and thus Internet wide communal cookie jar. But the instant we have individual per-site cookie jars, all cookie-based tracking disappears and you could say that tracking is terminated.

I'll repeat that it would be utterly wonderful to see this added into Chromium so that all of those non-Google users of the Chromium browser engine — Brave, Edge, Opera, Silk, Vivaldi and many others — could also terminate tracking. I doubt that this can be done with an add-on. It probably needs to be implemented somewhat deep in the browser core. Hopefully this notion will catch on.

https://blog.mozilla.org/security/2021/08/10/firefox-91-introduces-enhanced-cookie-clearing/

So what happened last Tuesday?

The Mozilla Security Blog posted an update on the privacy-enhancing changes they're continuing to make to Firefox. Sadly, they once again gave this new feature the rather milk toasty name "Enhanced Cookie Clearing." Really? They desperately need to find someone to snazz-up their naming department over there at Mozilla. Instead of the yawn-inducing "Enhanced Cookie Clearing" name, I was thinking of something more along the lines of how about "Website Historyectomy." That's catchy. *"With Firefox 91 you just press the "Website Historyectomy" button and Firefox immediately, completely and utterly forgets everything — and I mean everything — it ever knew about that website or that you ever visiting there."*

In any event, Mozilla explains it this way, and provides some additional detail:

*We are 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 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 Total Cookie Protection, Firefox 91's new approach 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". 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 browsing history or site-specific settings and permissions.*

*Firefox allows you to clear all cookies and other site data for individual websites. Data clearing can be used to hide your identity from a site by deleting all data that is accessible to the site. In addition, it can be used to wipe any trace of having ever visited the site from your browsing history.*

*To make matters more complicated, the websites that you visit can embed content, such as images, videos and scripts, from other websites. This "cross-site" content can also read and write cookies and other site data.*

*Let's say you have visited facebook.com, comfypants.com and mealkit.com. All of these sites store data in Firefox and leave traces on your computer. This data includes typical storage like cookies and localStorage, but also site settings and cached data, such as the HTTP cache. Additionally, comfypants.com and mealkit.com embed a like button from facebook.com.*

*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 storage for comfypants.com, Firefox deleted the storage of comfypants.com and left the storage of any sites embedded on it (facebook.com) behind. Keeping the embedded storage of facebook.com meant that it could identify and track you again the next time you visited comfypants.com.*

*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, if you clear site data for comfypants.com, the entire cookie jar is emptied, including any data facebook.com set while embedded in comfypants.com.*

*Now, if you click on Settings > Privacy and Security > Cookies and Site Data > Manage Data, Firefox no longer shows individual domains that store data. Instead, Firefox lists a cookie jar for each website you have visited. That means you can easily recognize and remove all data a website has stored on your computer, without having to worry about leftover data from third parties embedded in that website.*

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 Firefox 91 is extend this stovepiping isolation model beyond 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 1st party domain-based containment has been provided — and I hope that this is the future of web browser architecture — the only remaining trackable signals being sent by browsers are its static query headers, which are also in the process of being deliberately blurred, and 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. 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 user's explicit permission. All that was only added by techies because it's cool, not because it serves any clear purpose. 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 really any less of a privacy violation.

So, as regards cross-Internet tracking, the handwriting really does seem to be on the wall: News Flash!! Users don't like the idea of being tracked. Many might 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 tracking, then hell no! Are you nuts? And Google appears to be aware that their days of secretly tracking are numbered, too.

# Security News

**Facebook finally adds end-to-end encryption to Messenger**

It was back in March of 2019, more than two years ago, that Facebook's CEO Mark Zuckerberg proudly stated, as if they'd just discovered it, that <quote> *"the future of communication will increasingly shift to private, encrypted services where people can be confident what they say to each other stays secure and their messages and content won't stick around forever."* Many industry observers rolled their eyes at this, since even two years ago at this announcement, which was only to express an intention, Facebook was arriving quite 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, 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. Since 2016, we've offered the option to secure your one-on-one text chats with end-to-end encryption. In the past year, 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, if you choose.*

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 conversation 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 was, but it was at least 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 from everyone else's.

And, in case anyone listening to this podcast actually uses Facebook's Messenger system (I'm not judging mind you), you might be interested in knowing that they also announced that they had updated its expiring message feature within end-to-end encrypted chats. 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 5 seconds 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 DM's to 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 move forward toward this goal of mostly catching up with what other platforms and 3rd party solutions have long been providing. Better late than never. Eventually they might even turn it on by default.

**Exploitation of PrintNightmare has begun**

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 (CVE-2021-1675 / CVE-2021-34527) in Windows' print spooler service to spread laterally across a victim's network as part of a recent ransomware attack, according to Cisco Talos Incident Response research. 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 "Magniber" Ransomware Uses PrintNightmare**

Last Wednesday the CrowdStrike security blog was titled: "Teaching an Old Dog New Tricks: 2017 Magniber Ransomware Uses PrintNightmare Vulnerability to Infect Victims in South Korea"

> *CrowdStrike recently observed new activity related to a 2017 ransomware family, known as Magniber, using the PrintNighmare 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 PrintNighmare (CVE-2021-34527) vulnerability was disclosed, CrowdStrike intelligence assessed the vulnerability will likely be used by threat actors as it allowed for possible remote code execution (RCE) and local privilege escalation (LPE). This assessment proved accurate in light of the recent incident.*

**Crypto-mining botnet modifies CPU configurations to increase its mining power**

Ya gotta love this. A new crypto-mining botnet has begun modifying the operating configuration

of the CPU's on Linux servers it gains access to in order to increase the performance and efficiency of its cryptocurrency mining code.

Specifically, the mining code is modifies one of its processor's MSRs (Machine Specific Registers) to disable the CPU's hardware prefetch. Hardware prefetch is an optimization that's enabled by default because, as its name suggests, it enabled its processor to load data in its cache memory based on the operations that are likely to be required in the near future.

The security firm Uptycs spotted a crypto-mining botnet that was breaching Linux servers, downloading the Linux MSR driver, then using the Linux driver to disable hardware prefetching before installing a version of 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 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 prefetched data needs to be stored somewhere. So prefetching can be a bad thing of 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 caches.

So, in this case of an 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, but also because, now, a generally very useful speed optimization will have been disabled.

**NortonLifeLock and Avast are merging their users**
... in a deal valued between $8.1 and $8.6 billion.

To provide a bit of context and history, before this merger with Avast today, NotonLifeLock had acquired the German A/V maker Avira for $360 million in cash last December. And almost exactly one year before, NortonLifeLock was formed in November 2019 when Symantec sold its enterprise business to Broadcom for $10.7 billion, then rebranded their remaining consumer and small business operations 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 shares. The combined company will have dual headquarters, one in Prague, Czech Republic, and the other in Tempe, Arizona.

And most significant for us moving forward is that the new combined company will have a user base estimated at half a BILLION users with a "B". Avira is gone. Avast is gone. and though I've never really taken NortonLifeLock very seriously, they are clearly becoming a larger player in the online consumer and small business security scene.

**ASUS updates 207 motherboard BIOSes!**
https://www.asus.com/microsite/motherboard/ASUS-motherboards-Win11-ready/

As we know, the fact that Windows 11 would require support for version 2.0 of the Trusted Platform Module (TPM v2.0) came as surprising and unwelcome news to many Windows 10 users whose hardware lacks TPM v2.0. 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 v2.0, 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 Windows 11 easier and in some cases possible for users of their motherboards by offering updated BIOSes for 207 different motherboards.

Some AUSU BIOSes may already support Windows 11's requirement for TPM v2.0 and might only need to have one of two BIOS settings enabled depending upon the processor maker. That would be either: Intel Platform Trust Technology (Intel PTT) or AMD Platform Security Processor. Turn whichever one of those you see and if you have TPM v2.0 you'll be good to go.

But, since the BIOS might be somewhere some users fear to tread, an alternative is to simply download and run ASUS' new BIOS for your particular motherboard. It will load with TPM v2.0 deliberately pre-enabled and thus ready for the forthcoming upgrade to Windows 11.

I have a link in the show notes to ASUS' page, or I'm sure you can just find it at ASUS under getting ready for Windows 11.

# Errata

**L.T. Southall / @oiliswell**
"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. I do know what I'm talking about having spent 40 years in that business."

# Closing the Loop

**Darragh Duffy / @darraghduffy**

*Hi Steve, I just listened to episode 831. Another suggested reason why Apple are implementing a portion of CSAM on the device is that there's a suggestion that Apple is going to encrypt iCloud backups at some point in the future (currently they don't). Congrats in the 16 years of great work and security insights!*

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. We made the point 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 defacto 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, 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 their 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 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 a useful safety-value for law enforcement, allowing Apple to respond at least a bit and in specific circumstances. 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 tends to limit what can be done. Super-protecting photos of their cat and last night's lasagna is likely not high up on most people's list. And if it is, don't upload them to the cloud. These days, most people understand that once something is "in the cloud" it's never possible to really remove it.

**Joe Lyon / @Joe_Lyon_**

*The "tyranny of the default": In Germany there is about a 12% rate of organ donation. In neighboring Austria that rate is 99%. In Denmark the rate is 4% and in neighboring Sweden it is 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. In Sweden (where the rate is 89%), you check the box if you would like to opt OUT of the organ donor program.*

**Paul Durham / @PaulDurhamZA**

*Hi Steve, you and your incredible Security Now have inspired me to do something about the time gap between available & applied software & firmware updates, and how this gap leaves the consumer vulnerable. I am building a platform to allow developers & vendors to broadcast the availability of firmware & 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.*

**CatGuy / @Nomiswalcal**

*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 AP 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 2.4ghz Dlink). But it occurs to me to question if there is risk here too.*

---

# Microsoft's Culpable Negligence

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.

What follows is what I started off writing. So it starts off sounding like any other Patch Tuesday update...

**Patch Tuesday Thoughts**
Last Tuesday, Microsoft released fixes for 44 security vulnerabilities with 7 of the vulnerabilities being rated critical and three of those being 0-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 8 were information disclosures.

The affected Microsoft products included .NET Core & Visual Studio, ASP.NET Core & Visual Studio, Azure, Windows Update, Windows Print Spooler Components, Windows Media, Windows Defender, Remote Desktop Client, Microsoft Dynamics, Microsoft Edge (Chromium-based), 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 it's 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 RCE 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.

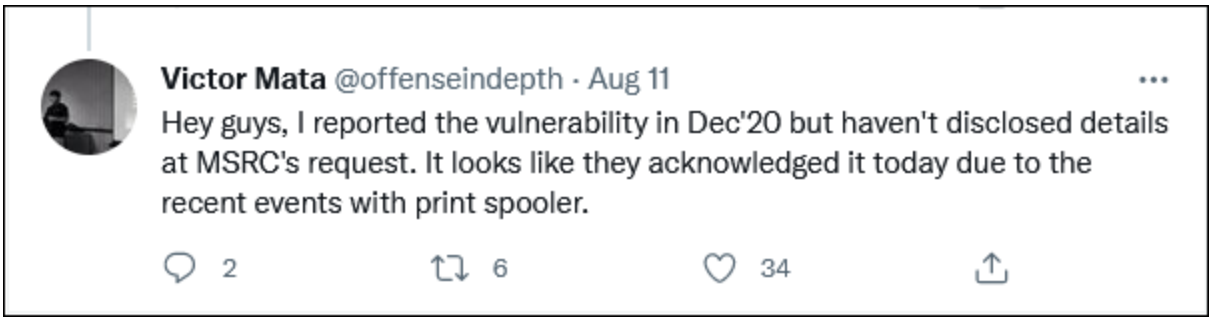


So now we're here again, with another newly disclosed Windows Print Spooler RCE. That's not good. It's really not good. But what's difficult to understand is that we're also told that Microsoft was first made aware of this problem way back in December of 2020 by Accenture Security's Victor Mata of FusionX. So... another remote code execution vulnerability in the Windows Print Spooler, which Microsoft has known about since December? And it's now 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 last Wednesday…

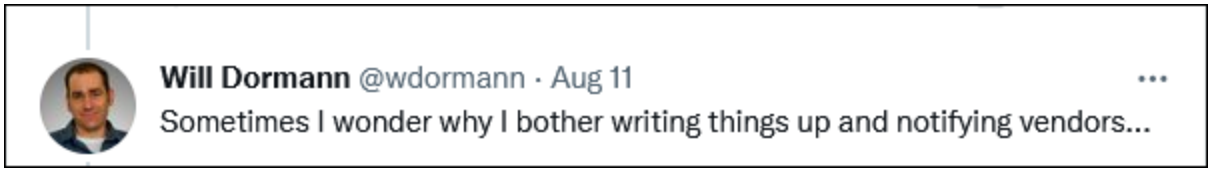


Yeah. I'd wonder, too.

Microsoft is nothing if not a savvy software publisher with effectively unlimited financial resources. Microsoft's current cash on hand is $130 Billion dollars. $130 Billion dollars of cash just lying around right now. They could have afforded to hire a talented coder to just 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... Shhhhh! 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. 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, are we? We've not blaming the faulty software which those attackers used to gain their foothold, exfiltrate their victim's proprietary data and encrypt their victim's machines. We're blaming the attackers now. 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 better stop attacking us, or else!" ... while Microsoft sits on another Remote Code Execution flaw in Windows Print Spooler for eight months.

Microsoft is 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 dollars 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 the 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. 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 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?

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 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.