BlueKeep and DoH

Description: This week we examine a widespread Windows breakage introduced by last month's Patch Tuesday. We look at several things Google changed in their just-released Chrome 78, news from the Edge, the status of attacks on Intel chips, a new attack on publicly exposed QNAP NAS devices, the significant risk of trusting managed service providers, the downside of apps for autos, and worries over Chinese-made drones. We then finish by coming back to look at news on two other fronts: the escalating controversy over DNS-over-HTTPS (DoH) and the commencement of the long-awaited BlueKeep vulnerability attacks.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-739.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-739-lq.mp3

SHOW TEASE: It's time for Security Now! with Steve Gibson. I'm Jason Howell, filling in for Leo once again. We've got news from the Edge, Microsoft Edge; an attack on Internet-connected QNAP NAS devices; why connecting to rental cars via apps might not be such a great idea after all; Steve checks in on the DNS-over-HTTPS controversy; and the BlueKeep attacks have begun. Steve Gibson's going to break it all down for you next on Security Now!.

JASON HOWELL: This is Security Now! with Steve Gibson, Episode 739, recorded Tuesday, November 5th, 2019: DoH and BlueKeep.

It's time for Security Now! with Steve Gibson and myself, Jason Howell, filling in for Leo, week three. How you doing, Steve?

Steve Gibson: Great. Good to be with you for - this is three of our four; right?

JASON: That's right.

Steve: We get you one more week, next week, and then - actually, and then I'm up there on Saturday, November 30th.

JASON: That is it, indeed.

Steve: The fires are out, and things are calming down now. The air quality has been restored.

JASON: It's really interesting dealing with that scenario because there for, like, four or five, six days, it was everything that anyone was talking about, you know, locally. It was just like fire was on the mind. And then the second it dissipates it's like, quite literally, back to normal life.

Steve: That's right.

JASON: It's just the kind of counterpoint between one extreme to the other is very interesting to me. But yeah, we are here and ready to talk some security. We actually have some folks sitting in the audience here today.

Steve: We have a guest audience.

JASON: Yes, we do.

Steve: A packed house.

JASON: Jim Willard and Tim Willard, both from Michigan, and Sharif Vallejo. Where are you from, Sharif? Oh, sorry, I thought that was your last name. I was like, Sharif Vallejo, where are you from? Sharif is from Vallejo. So there we go. He's here from Vallejo. So we've got three guests sitting in for the show today. So Steve, you've got to talk some serious security.

Steve: Very cool.

JASON: It's got to be serious today.

Steve: So we have Episode 739 for November 5th. And there were two things I want to sort of update us on. One is DoH, and the other is BlueKeep. DoH, of course, is DNS-over-HTTPS, which has turned out to be surprisingly controversial for everyone except users, who think it's a good idea. But, oh, that's really causing some fur to fly. So we're going to talk about that. And then of course BlueKeep is the much-anticipated end of the Internet as we know it, which everyone keeps anticipating. Well, finally we have attacks, but they are not what the Internet expected. But they are what we were expecting on the podcast. So we're going to talk about that. We're going to end the show talking about that.

But there's a whole bunch of interesting stuff happened this week. We've got a - I was affected by this - a widespread Windows breakage which was introduced by last month's Patch Tuesday. It had been bothering me ever since the second Tuesday of October. And in putting the news together for this show, I stumbled upon its explanation. I implemented it, and it fixed the problem. So I'm sure I'm not alone, and I wouldn't be surprised - let's see. This is the first Tuesday of November. I wouldn't be surprised if Microsoft unfixes or unbreaks what they broke next week. And let's hope so because this has to be affecting lots of other people. But we have a fix for it right now.

We also have, speaking of fixing what was just broken, Chrome 78 has made some important changes to what they did in Chrome 77. So we'll discuss those. We also have our new segment, News from the Edge, which is to say about Microsoft's Edge browser. It's a little less dramatic than the name.

JASON: I like that. It's catchy.

Steve: News from the Edge. We've also got the status, nearly two weeks downstream, the status of the attacks on Intel chips which were the whole Spectre and Meltdown thing. And, you know, it does make you feel a little bit old when you realize that, okay, wait, Spectre and Meltdown, that was the beginning of 2018. That's, like, wow. Where did these two years go?

JASON: It's kind of crazy.

Steve: It doesn't seem like it was that long ago.

JASON: Unh-unh.

Steve: We also have an important new attack on publicly exposed QNAP Network Attached Storage boxes. I know that our listeners are bullish on those because after I did my conversation about - we called it "The Joy of Sync" podcast a couple months ago, where I talked about my search for a really good way of synchronizing multiple locations and keeping directories synchronized. QNAP has some of their own software. And so I saw a lot of tweets from people who are using QNAP NAS devices. They need to make sure that they update to the firmware that just came out on November 1st. So we'll talk about that.

We also have an interesting report about the significant risk of trusting managed service providers and maybe the need to push back on the kind of access that they by default want to have into their clients' networks. We've got the downside of apps for autos, and worries over Chinese-made drones. So a whole bunch of fun stuff to talk about this week. It's going to be another great podcast as we approach number 999, which of course we all know is the end of Security Now!. And with that will be the end of all security problems on the Internet. So, yeah.

JASON: This is breaking news.

Steve: Well, I only have three digits, Jason.

JASON: Well, it's breaking news to me that Episode 999 is the end of security issues.

Steve: My whole system is based on three digits. And so it's like the Y2K of Security Now!, and it's lights out. And Leo's been thinking, you know, that's about, okay, that's how long ago? When's that going to be? Okay, that kind of, you know, probably just works out. He'll be off on a cruise when we wrap around from 999 to 000. Actually, funny story.

JASON: I can even clean up the mess while he's gone. It's fine.

Steve: I was eating at - I had a favorite Chinese restaurant which had a computer-based system for managing, you know, that they put all their orders into and everything. And this was a Y2K problem because I'm sure they never licensed the software. And so when I first went in after the New Year, after Y2K, their systems had gone from 1999 to 19100. So literally, so the 19 was hard-coded into the software, and the 99 went to 100. So the checks came out, and they were five digits. It's like, 19100. So that was a bit of a problem.

JASON: Whatever works for you, business owner. Okay.

Steve: Yeah. Anyway, they knew I was a computer guy. They said, "What should we do?" And I said, "Well, you know like how Leap Year works; right?" And they go, "What?" I said, "Okay, just go back, set everything back four years. The year won't be right, but at least things will kind of line up correctly."

JASON: Yes.

Steve: And anyway, so that's how they operated. Yeah, 19100.

JASON: And I'm curious to know if they're still dealing with that every four years.

Steve: Actually, they went out of business the summer of 9/11, the attacks of 9/11, because I don't know if you remember, but the restaurant business took a big hit...

JASON: Yeah, it did, I remember that.

Steve: ...after the 9/11 attacks because everyone was a little freaked out and stayed home and ate popcorn and watched TV to see if anything was going to happen next. So anyway, that kind of did them in, even though they were the best Chinese restaurant place in Southern California.

JASON: Well, that's unfortunate.

Steve: And in other trivia - no.

JASON: Yeah, I know. Where are we going to go next? All right, so we're starting with, I guess - I guess this is the Picture of the Day, kind of? Sort of?

Steve: Well, it is the Picture of the Week because it ties into our first story. This was literally what I have been facing ever since my Windows 7 machine updated after the second Tuesday of October. And there it shows me trying to bring up the Digikey.com site. Digi-Key is one of my favorite and is actually my go-to site for electronic components, Mouser being second choice. And here it shows down in the lower left, "Establishing secure connection." And that's where it gets stuck. And that's trying, after I hit Refresh, after we got the unhappy web page icon, and it says "This site can't be reached; www.digikey.com took too long to respond." And then it gives you some generic problems.

So Windows broke something with last month's Patch Tuesday, which Microsoft has now acknowledged. And my sense is that, since this is a core protocol thing, it can't be just me trying to get to one particular website that's been having problems. So as we've been recently discussing, I often, like you, Jason, well, you're a stronger user of Chrome. I'm fickle. So I will be...

JASON: You're a flip-flopper.

Steve: I'm a flip-flopper.

JASON: Okay.

Steve: Exactly. I'll often be using Firefox. Normally I actually have Firefox open statically in my lower left-hand display. That's just kind of where it lives. Over on the right is Windows Explorer. In the center is what I'm doing. So I sort of have, like, reference stuff around me. And then I have sort of a fixed layout for where things go, like what goes on which screen. So it's been my habit to open up a Chrome session in the center screen when I want to browse around Digi-Key for some random electronics gizmo. Except that Chrome stopped being able to display Digi-Key's site. It would, you know, the thing would spin and eventually complain that it was unable to obtain a secure connection, or any at all.

And, I mean, this has really been an issue. In the past few weeks I've hit F12 to open Chrome's developer window, to examine the exact error code that was returned. I thought that perhaps Digi-Key might have somehow misconfigured their site. Because, I mean, again, I've used these guys for years, and never have I had a problem. So I remember I went over a couple times actually to Ivan Ristic's SSL Labs site and had them scan Digi-Key. But it got top marks from Ivan. So it didn't seem to be something there.

And one of the more curious things was that Firefox never had any problem. While Chrome wouldn't open Digi-Key, Firefox - it just popped right up in Firefox. But I wanted to use Chrome, and I couldn't. So it turns out it wasn't just with Digi-Key. Under Chrome over the past few weeks other random sites all over the Internet started having trouble. And since I learned a long time ago that things that I'm experiencing are likely to be widespread, since I'm not doing anything weird, on a number of occasions over the last nearly a month I would Google "Chrome connection errors" or "Chrome connection timeout."

I mean, I figured other people would be, like, noticing this and having problems, thinking that Google must have broken something on Chrome that would therefore be affecting lots of people, very much like they did break something with Chrome 77. Well, not actually they broke it, and we'll be talking about this in a minute, where we talked about how, if you had the un-updated version of Symantec's Endpoint Protection, people were getting the "Aw, Snap" error from Chrome.

But anyway, so as I said, as I was doing the research for today's news, I stumbled upon the cause. And I'm telling everyone who's listening in case anyone else has been, like, pulling their hair out with the same sort of problem, not being able to use Chrome in certain instances.

So here's the story. I don't know how Microsoft managed to do this, but they broke some, well, they added support that is enforcement for an enhanced anti-man-in-the-middle protection involving TLS handshakes, which immediately broke Windows' ability to connect as a client or as a server to a significant number of the Internet's websites. And what was interesting, when I stumbled upon that, was it was like this "aha" because it perfectly explained why Firefox has continued to work because, as we know, Firefox does not use Windows' underlying security stack. It brings along its own. But Chrome does run on top of the Windows security protocol. So if Microsoft had done something to Windows, Chrome would be affected; Firefox would not be.

So once I understood what was going on, I immediately disabled what had proven itself to be a not-yet-ready-for-primetime protocol handshake feature, rebooted my machine - though some reports suggest it isn't necessary to reboot - and Digi-Key's site popped up on Chrome with no problem. So BleepingComputer - which is where I finally, it's the only place I saw this noted, like in reading everything over the last couple weeks, so hats off to them.

They said in their note: "Microsoft has acknowledged a new issue affecting several Windows versions" - which is, yes, everything currently being updated, meaning 7 through 10 inclusive, that Microsoft says, or BleepingComputer reports of Microsoft - "could lead to Transport Layer Security and Secure Sockets Layer connections intermittently failing or getting timed out." And in my experience, never getting made in the first place.

They said: "This bug is caused by the security-related enforcement for the" - and then we have a CVE, it's 2019-1318 - "TLS spoofing vulnerability, which leads to Windows devices experiencing failures and timeouts during TLS DHE dot star," essentially. That's all of the TLS cipher suites involving ephemeral Diffie-Hellman, which is what the DHE, Diffie-Hellman Ephemeral, DHE stands for. "This happens only when the devices are trying to make TLS connections to devices without support for the Extended Master Secret (EMS) extension."

So what Microsoft wrote was: "Connections between two devices running any supported version of Windows should not have this issue when fully updated." In other words, says Microsoft, as long as you have Windows, fully updated Windows at each end, at the client and the server, well, what's the problem?

JASON: That'll solve the problem for everyone, won't it.

Steve: Earth to Microsoft. News flash: Not everyone is using your Windows operating systems at each end of the connection on the Internet, especially not on the server end, where IIS has been steadily losing steam and market share to Nginx, Cloudflare, and other providers through the recent years. So Microsoft's notice of this is headlined: "Transport Layer Security (TLS) connections might intermittently fail or time out when connecting." Yeah.

And so I've got a link in the show notes to their notice. Fortunately, there is a registry tweak which can end this trouble by redisabling Windows' new and apparently misguided enforcement of this Extended Master Secret handshake. I have a picture that I took of my screenshot of Registry Editor, where I added two keys. You add two keys under - and I'm not going to try to go over in detail, but they're easy to add for anyone who's comfortable playing with the registry, as most veteran Windows users have had to learn to be. It's under Computer, HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control, and then SecurityProviders. Under SecurityProviders is SCHANNEL in all caps. And there you add two DWORDs, two 32-bit DWORDs for DisableClientExtendedMasterSecret and DisableServerExtendedMasterSecret. And you set those DWORDs to "1," and that turns off Windows' failing attempts to negotiate an Extended Master Secret with the other end that doesn't know about that. And, incredibly, Windows refuses the connection if it cannot obtain this Extended Master Secret negotiation.

So I just wanted to give everyone a heads-up. I mean, my pain is over, thanks to having stumbled upon this. So a huge and profound thank you to BleepingComputer for pointing this out. And who knows how long I and doubtless many others would have been stuck without Chrome working on sites that would fail what is now an enforced handshake over on the client side, and presumably over on the server side, since there are options to disable either the initiator of the connection, thus the client, or the acceptor of the connection initiated by a client, so on the server side.

So again, it must be the case that Microsoft realizes, whoops, this is not what we wanted to have happen. And I wouldn't be at all surprised if next Tuesday we don't have this problem resolved at their end. But if anybody else has been beset by this, I mean, so the protection would be nice to have. This does solve a connection interception man-in-the-middle spoofing problem which TLS can have with the ephemeral Diffie-Hellman key agreement negotiation. So it would be nice to have it. But despite the fact that this has been in the spec for at least four years, I saw references to 2015 around, I didn't bother to go into big detail, but it's one of those things where, yeah, it'd be nice to have it. But it sure needs to fail gracefully. And again, I had a lot of stuff to cover for the podcast, as was obvious at the beginning of the show, so I didn't spend time digging into this.

One of the problems with man-in-the-middle spoofing attack protections is they tend to be subject to downgrade attacks, meaning that the man in the middle can, since they're able to intercept the connection as part of the attack anyway, they're able to say, oh, we don't support Extended Master Secret. Sorry. Which is to say that, if the endpoints allowed Extended Master Secret to be soft supported, which is to say only supported when each endpoint agreed, then the first thing a man-in-the-middle attacker would do is claim not to support it, thus disabling it. So it does make sense that Microsoft wants to force its enforcement, but it's breaking the Internet. So unfortunately, we're not ready for, obviously, for its enforcement.

So I don't know what's going to happen, you know, because Microsoft tried to bring it up, and it broke a bunch of things for me. And I can't be alone in this. So as I said, it'll be interesting to see what happens next week. And Microsoft knows they did this. Their advice is, oh, well, update Windows at each end. It's like, no. That's just not practical.

JASON: Yeah, not always an option.

Steve: Hello. And less of an option moving forward as the world leaves IIS as the web server platform of choice. So it doesn't look like that's going to happen any time soon. And I was curious, too, because I did go back to SSL Labs. I remembered that they showed what the server that Digi-Key had identified itself as using was, and now I don't remember. But it was some cloud provider web hosting platform thing. So it wasn't like Nginx. It was something that maybe there isn't that large a market share and so not that many servers are now in a position where they don't support it. But I couldn't get to a website that I want to get to using Chrome. So that's not okay. So anyway, I would imagine this will be of interest to some of our listeners. So, hope so.

And speaking of Chrome and 77 and 78, 78 is just out, and several things have been fixed. One is that the thing they did last week with Chrome 77, which turned out to have a problem with Symantec, which was because they were enforcing Microsoft's Windows Defender Code Integrity Checking, which would prevent things not signed by Microsoft from being loaded into the Chrome process, the Chrome renderer. That turned out to break systems that were using Symantec's Endpoint Protection. But it turned out that there were more things that it broke, not just Endpoint Protection, but also users were getting the "Aw, Snap" crashes if they used PC Matic, which is that cheesy, advertised on late night TV AV. But also Palo Alto Networks traps security products, as well as something known as the Print Audit Infinite tool, which I've never heard of, but apparently it's used for tracking document printing on LANs.

So it wasn't just Symantec. There were other apps which were also broken by this. Which is to say, for some reason, these things were injecting their code into Chrome's process. And this Code Integrity Checking allows Chrome to defend itself against that code injection. And that's good for everybody. We want that to be on. But Google has decided to turn it off. And so with Chrome 78, it is off. They have said, however, that they intend to turn it back on in the middle of this month. So those companies that have been alerted to the fact that their solutions are not compatible with Code Integrity Checking need to figure out how to fix that. Apparently it's possible because we know that, if you update to the newer version, the current version of Symantec's Endpoint Protection, then this doesn't cause that problem.

So I guess this was probably necessary to clearly inform the companies that needed to that they've got to get their act together, or their users are going to be uninstalling their programs from their computers which are causing Chrome to break. And I would imagine that Google would win this battle of incompatibility, rather than the Print Audit Infinite tool and PC Matic. Anyway, so that has happened.

Also, last Thursday, on Halloween, which was when Google released Chrome, it's 78.0.3904.87. And when I saw that I checked, and that's the one I had. Not only to backpedal on, as I said, that enforcement of this Code Integrity Checking, but every bit as urgently to patch a new Chrome zero-day which they had just been informed of by Kaspersky. And it was being exploited in the wild. I have a link to the googleblog.com information about this.

They said: "The stable channel has just been updated to 78.0.3904.87 for Windows, Mac, and Linux, which will roll out over the coming days and weeks." Hopefully days. They said: "Note: Access to bug details and links may be kept restricted until a majority of users are updated with a fix. We will also retain restrictions if the bug exists in a third-party library that other projects similarly depend on, but haven't yet fixed. This update includes two security fixes." They said: "Below, we highlight fixes that were contributed by external researchers. Please see the Chrome Security Page for more information."

So there are two updates. The second one was the one that was sent to them by the guys at Kaspersky Labs. And to Google's credit, they received the information on Halloween, and they had 78 pushed out, or available in the stable channel, that same day. So they understood this was important, and they fixed it immediately. Kaspersky's disclosure, although Kaspersky pulls back, of course, from describing the details, Kaspersky's disclosure did have some wonderfully juicy details.

They said: "Kaspersky's Exploit Prevention is a component part of Kaspersky products that has successfully detected a number of zero-day attacks in the past," and of course we've talked about those. "Recently, it caught a new unknown exploit for Google's Chrome browser." They said: "We promptly reported this to the Google Chrome security team. After reviewing the proof of concept we provided, Google confirmed there was a zero-day vulnerability and assigned it CVE-2019-13720. Google has released Chrome version 78.0.3904.87 for Windows, Mac, and Linux; and we recommend all Chrome users update to this latest version as soon as possible."

And as it happens, I had fired up Chrome at the beginning of the work on the podcast yesterday. And it wasn't till I went to Help About that that gave Chrome a little kick in the pants to update itself. I was on 77 still, despite having just launched it. So it's worth making sure the next time you launch Chrome that you've got 78. So at this point it's very targeted. It's unlikely that people who are not - I'm trying to think where it was. Anyway, I'll get to it in a second. But it looked like it was only a watering hole attack affecting a limited number of people. But it was an unknown zero-day at the time, so nice that Kaspersky caught it.

They said: "We're calling these attacks Operation WizardOpium." They said: "So far we've been unable to establish a definitive link with any known threat actors. There are certain very weak code similarities with Lazarus attacks, although these could very well be a false flag. The profile of the targeted website is more in line with earlier DarkHotel attacks that have recently deployed similar false flag attacks." So these guys are at the top of their game. They said: "More details about this and recent DarkHotel false flag attacks are available to customers of Kaspersky Intelligence Reporting. For more information, contact intelreports@kaspersky.com."

But the tech details, they said: "The attack leverages a watering hole-style injection on a Korean language news portal." Okay, that's what I meant when I said rather targeted. Unless you happen to be visiting Korean language news portals - now, we don't know that's the only place this is being used. This is where they found it. They said: "A malicious JavaScript code was inserted in the main page, which in turn loads a profiling script from a remote site. The main index page hosted a small JavaScript tag that loaded a remote script from" - and then it's a code.jquery.cdn.behindcorona.com site. "The script then loads another script named dot charlie dot," and then they redacted the dot js.

"This JavaScript checks if the victim's system can be infected by performing a comparison with the browser's user agent, which should run on any 64-bit version of Windows and not be a WOW64 process," meaning a 32-bit process running on a 64-bit OS. "It also tries to get the browser's name and version. The vulnerability tries to exploit the bug in Google Chrome browser, and the script checks if the version is greater or equal to 65." The current Chrome version, of course, as we know, is now 78, which is immune to this.

"If the browser version checks out, the script starts performing a number of AJAX requests to the browser's controlled server at behindcorona.com, where a path name points to the argument that is passed to the script xxxxx.php." Again, they redacted. "The first request is necessary to obtain some information important for further use. This information includes several hex-encoded strings that tell the script how many chunks of

the actual exploit code should be downloaded from the server, as well as a URL to the image file that embeds a key for the final payload and the RC4 key to decrypt these chunks of exploit code. After downloading all the chunks, the RC4 script decrypts and concatenates all the parts together, which gives the attacker a new JavaScript code containing the full browser exploit."

So these guys went through, I mean, they jumped through serious hoops in order to obscure and encrypt and protect from observation the JavaScript code for the exploit. So remember, this thing is checking from Chrome 65, which presumably is where this problem began, anything greater than 65 at the time. So that suggests this thing has been around and has been effective for some period of time.

"The analysis," they wrote, "we have provided here is deliberately brief due to vulnerability disclosure principles." In other words, they don't want to give this away while browsers are still in the wild that could be affected by this. They said: "The exploit used a race condition bug between two threads due to missing proper synchronization between them. It gives an attacker a use-after-free condition that is very dangerous because it can lead to remote execution scenarios, which is exactly what happens in our case," Kaspersky wrote.

"The exploit first tries to trigger the use-after-free to perform an information leak about important 64-bit addresses." So that's probably, as we know, address space layout randomization deobfuscation. They said: "This results in a few things. If an address is leaked successfully, it means the exploit is working correctly. A leaked address is used to know where the heap and stack is located, and that defeats the address space layout randomization" - yup, ASLR - "technique. And, three, a few other useful pointers for further exploitation are then locatable by searching near this address.

"After that it tries to create a bunch of large objects using a recursive function. This is done to make some deterministic heap layout" - we've seen that before, heap grooming, called "grooming the heap" - "which is important for successful exploitation. At the same time, it attempts to utilize a heap spray technique that tries to reuse the same pointer that was freed earlier in the use-after-free part. This trick could be used to cause confusion and give the attacker the ability to operate on two different objects from a JavaScript code perspective, though in reality they're located in the same memory region.

"The exploit attempts to perform numerous operations to allocate and free memory along with other techniques that eventually give the attackers an arbitrary read/write primitive. This is used to craft a special object that can be used with WebAssembly and FileReader together to perform code execution for the embedded shellcode payload. After decryption, the malware module is dropped as updata.exe to disk and executed. For persistence, the malware installs tasks in Windows Task Scheduler." That's a common means of getting something to run when the system boots or the user logs on.

So in short, from a user's standpoint, with Chrome before 78, and presumably back to 65, if a person were to visit a website, and without them clicking, doing anything, or knowing anything, their Chrome browser would download, install, and run a native application on that Windows machine, which is to say, this is as bad as it gets. We don't know how long this has been in the wild. We know apparently it's been effective since Chrome 65. And it's clear, I wanted to give our listeners a good sense for how much industry went into the creation of this. This was an expensive exploit to deploy. It is effective against the world's number one browser on the Internet, which we now know Chrome is. And it's very clear that anyone going to that particular news portal with Chrome of the appropriate versions, and Chrome is updating itself all the time so nobody would have had an older one, and that is until they update to 78, would get themselves infected.

So bravo for Google fixing this immediately. Thread synchronization problems are among the most difficult problems to fix. All the code is working. Everything's fine. The fact that they were blasting the heap and stack with allocations suggests that they were trying to obtain a pointer that would function with JavaScript still able to access the region of memory. So they would have a means then of writing to memory through JavaScript, and then getting execution of that memory through essentially an outside of JavaScript pointer.

So, wow. Again, as I said, this kind of synchronization problem, they are among the most difficult problems to find. They are very subtle. So first of all, somebody had to find this. Then they clearly put together a heavy-duty campaign in order to leverage this into remote code execution and then set this up on some site. You can use your imagination. Who would want to be installing malware into visitors of Korean news portals? So, wow. That's where we are today.

JASON: But at least you convinced me to use Firefox, so it doesn't matter anymore; right?

Steve: Very good point, yes. And maybe some of us haven't been able to use Chrome, thanks to what Microsoft did to Windows last month.

JASON: Right, exactly, all the signs are pointing away.

Steve: That's right.

JASON: All right. It's time we step over the Edge into this next segment.

Steve: So just a note to our listeners. This is not a big piece of news, but I wanted to put it on everyone's radar. The preview build for the new Chromium-based Microsoft Edge Stable version has been released on the Edge Insider site. This is the preview of the Edge browser that will be made widely available in January. Yesterday, as part of the kickoff at the start of Microsoft's five-day Ignite 2019 conference, the announcement was made that Microsoft Edge will become generally available and fully released, that is to say, the Chromium-based version, middle of January, on January 15th.

And so what's interesting is with this preview release, Microsoft Edge Stable will be at v79.0.309.7, which actually puts it one major version ahead of where Google's Chrome current stable version is. As we know, it's at 78. But on the other hand, Chrome will be past that point, especially if they rev Chrome in the middle of this month in order to turn Code Integrity Checking back on for probably Chrome 79. So we expect them to do that.

So from an academic curiosity standpoint, it's going to be very interesting, I think, Jason, to see how the world settles out with a version of Edge that's internally identical to Chrome. You know, same, I mean, they'll both be Chromium. They'll be from different companies. They'll have different UI Chrome. But so will Edge take some of Chrome's market share? Will anyone know or care what's under the hood of their browser? I don't think they do now.

JASON: I don't think they do now, yeah.

Steve: No. So, you know, and so I guess that makes me wonder what's been moving people away from Edge in the first place. I have no idea. But it's going to be interesting to see how this develops. In general, you know, I applaud this move by Microsoft. It must have been so difficult to discard all of the work that was done on their brand new, very own new browser that replaced the old creaky IE engine. That was the ChakraCore. But this clearly was the sanest choice. Why run independent projects for something that

has become a commodity? Just, I mean, this makes sense. But do you think that people using Windows 10 with Edge, I mean, they must be using Chrome. I mean, they must be downloading Chrome and choosing to use it instead of Edge for some reason.

JASON: I mean, yeah, maybe that's just a force of habit. If they already happen to have all their Google accounts, Chrome makes it really, really appealing from the perspective of syncing all of your information, having that be another data point into your Google account. So, yeah, I don't know, either. I'm also admittedly not really much of a Windows user. So I don't know through experience what that would be like. I know that if I ended up on a Windows browser, like my habit would be to download the Chrome browser because it logs into the Google accounts and because that's just where I have all my data flowing through anyway.

Steve: Yeah. I think that must be it. And of course, if you are a user of Google Search, Google is actively promoting Chrome. It's like, are you ready to switch? Why haven't you switched yet? Let's switch you right now. Wouldn't switching be a good idea? It's faster. It's more stable. It's Google.

JASON: It's what you want.

Steve: And so I'm sure that people are like, oh, yeah, I don't want to use this stinky old Microsoft thing that works just fine. I want the shiny Google version. So, yeah.

JASON: Right, right.

Steve: So we are two years downstream from our first encounter with microarchitectural data sampling vulnerabilities. That's the fancy rollup catchall classification name that groups together all of those side channel attacks on Intel and, to a lesser, but still significant degree, AMD and other processors, which as we know started coming to light right from the start of 2018. I mean, it was podcast one of January 2018 that was about Spectre and Meltdown, which puts it nearly two years ago. We had, as we know, we first had Meltdown, and then Spectre. And then they were followed by a few variants and extensions of them. Then we were hit with PortSmash, then ZombieLoad, then RIDL, and then more recently Fallout. So there have been a bunch. And of course these are distinctive from the DRAM hammering attacks which have been demonstrated to be effective in the wild.

So here we are, though, nearly two years downstream of all that, and one glaring aspect remains true, which is despite there never having ever been any sign of any of these theoretical processor failings being leveraged into a practical attack, we have all nevertheless sacrificed around 20% of our previous system performance in order to prevent these attacks that have never materialized. So of course we can't prove a negative. There's no way to know whether if the industry had not responded as affirmatively as it collectively did, we would have eventually discovered these problems in the wild. And we'll likely never know. Hopefully Intel has learned a lot from all of this.

What they've done so far and what we've done, you know, what they've done with microcode patches and what the OS industry has done by implementing kernel-level embracement of those patches is to, as has been noted by people who have been doing benchmarks, it's now felt that we're paying about a 20% performance price, a penalty for this, just because we've turned off things that were all about optimizations at the microarchitecture, which it turned out were not safe.

So at this point what we know is that Intel now has a profound understanding of these problems. And remember, there was some mention of this theoretical possibility early on,

which people just sort of ignored because we were ignoring all kinds of things in the early days, which we now soberly understand we cannot afford to ignore any longer.

So Intel knows what the problem is. We know that their processor design pipeline is very long and deep, meaning which is to say that they start on a design years before it is in consumer silicon. So it's certainly the case that they immediately scrapped the designs as far back as they could, or as close as they could, and are bringing a whole new architecture out. But it's going to be years until we see those.

Probably at that time we will get everything we want. We will recover the performance that we've lost. We will have designs which are now fully aware of these problems and have new workarounds and probably a whole bunch of new patents filed by Intel and other processor manufacturers about how to not be susceptible to these microarchitectural data sampling problems while still developing for us the kind of performance that we were liking to become accustomed to. With any luck, we'll have some new chips in a few years that are fast and safe, both.

Okay. So QNAP. As I mentioned at the top of the show, I know from many instances of feedback I've seen after the podcast about the file synchronization solutions that I found, that many of our listeners are very happy users of their QNAP brand of network-attached storage devices. So I wanted to make sure that they and everyone were aware of the fact that many thousands of QNAP NAS devices are currently infected with malware which has been named QSnatch. Over 7,000 infections have been reported in Germany alone, and the malware is still spreading. Thousands more, many thousands more, are believed to be infected worldwide in what appears to be an ongoing outbreak.

And thank you. On the screen right now is a map of known QNAP NAS device infections. Clearly Germany is just buried. But so are the coasts. Doesn't look like there are many QNAP NASes in the more rustic regions of the U.S. But, boy, in the Pacific Northwest and in Southern California, California in general, it's a lot of problems.

JASON: Blanketed. Saturated.

Steve: Blanketed, thank you. Perfect word. And also on the Northeast is also just you can't see the country underneath that cloud cover.

So information on how QSnatch works is still scant. The only report comes from the National Cyber Security Centre in Finland, which was the first cybersecurity organization to spot the malware just last week. NCSC-FI, that's the Finland members, have not yet discovered how this new threat spreads and infects QNAP NAS systems. However, once it gains access to a device, QSnatch burrows into the firmware to gain reboot persistence.

So an analysis of the malware's code revealed the following capabilities. We know that it modifies OS timed jobs and scripts, so cronjob and init scripts. It prevents future firmware updates by overwriting update source URLs. So it prevents the machine from updating itself. It prevents the native QNAP MalwareRemover App from running. It extracts and steals usernames and passwords for all NAS users.

So these features, while describing the malware's capabilities, don't reveal its intent. So it's unclear if QSnatch was developed to carry out DDoS attacks, to perform cryptocurrency mining, or maybe just as a way to backdoor QNAP devices to steal sensitive files or host malware payloads for future operations. We don't know yet. One theory is that QSnatch operators are currently in the phase where they're building their botnet and will deploy other modules to it in the future. The analysts confirmed that QSnatch has the ability to connect to a remote command-and-control server, to download and then run additional modules.

So this is definitely something you want out of your QNAP NAS, if you're unlucky enough to have had it infected. And based on the patterns we're seeing, it looks like Europe is buried, and so is the U.S. wherever these things are. So for the time being, the only confirmed method of removing QSnatch has been performing a full factory reset of the NAS device. After performing a factory reset, users should install the latest QNAP NAS firmware update, which was last Friday, November 1st. QNAP released a firmware update which incorporates specific QSnatch protections. So QNAP themselves know what the problem is. They've got a firmware update.

I don't own a QNAP device. I bought a Drobo device back when Drobo was a sponsor of the show, and I have them at two locations, as our listeners know. So I can't speak to whether the NASes would automatically update, or could, if they haven't yet been infected. What we do know is that infections are aplenty, as that map showed. And once infected, your QNAP will no longer update itself. So I would recommend a manual factory restore and then update. Because anybody who has a QNAP device will want to be current, and you're going to have to make sure, apparently by manual means, that you get any existing infection scraped out of your machine.

Once you've done that, you'll want to change all passwords for all accounts on the device, since they may have been compromised; remove any unknown user accounts from the device; make sure that the firmware, of course, is up to date, and also that your apps are up to date; remove anything you don't recognize, any unknown or unused apps. Just as general proper behavior you would want to do that. Also install the QNAP MalwareRemover application through the App Center functionality. And, finally, set an access control list for the device so that you've locked this thing down.

We'll note that QSnatch is the fourth malware strain spotted this year that has targeted NAS devices. There was a ransomware strain that was impacting Synology devices, and there were two previous ransomware strains that were infecting QNAP devices. So QNAP has had two previous encounters this year. It really is clear, if you've got your QNAP device exposed to the Internet. And I should have started by mentioning that also. I mean, from what our listeners were saying, it looks like they appreciate the functionality of having their NAS, their QNAP NAS on the 'Net. That is, with a public-facing interface that allows them to access it remotely. That's the danger, of course.

If your QNAP is only on your LAN and is not exposed to the public Internet, then you never had a problem. You know, you're not one of those glowing orange circles on the map that we showed. But for what it's worth, I know that a lot of our users are using these things. Just take some time as soon as you can to make sure that you have not been infected, and do make sure that you update the latest firmware because this thing looks like it's serious, and we don't know what it's going to do next.

JASON: Yeah, sounds ugly.

Steve: Yeah. I mentioned at the top of the show just sort of a - I wanted to put this on people's radar. And that is we now have a report from Armor, the folks that have been watching and really doing a lot of forensics work on ransomware delivery. They now have positively identified at least 13 Managed Service Providers who've been positively identified as being the means for having pushed ransomware out to their client companies. The Managed Service Provider functions as essentially a large access hub for the ransomware. And once hackers compromise an MSP's network, they then use its remote access tools to deploy ransomware to often hundreds of companies, and that means thousands of computers.

I have a link to the report in the show notes for anyone who's interested in the details. And I'm not going to go through an enumeration of this. But we've referred to Armor in the past. This report adds six new identified managed service providers and cloud-based

service providers to the existing list that they had already known of, which brings the total of publicly identified MSPs and cloud-based service providers of just during 2019 to a total of 13.

Chris Hinkley, the head of Armor's Threat Resistance Unit, said: "This uptick in successful ransomware attacks against MSPs and Cloud-Based Service Providers," he said, "is a harsh reminder that organizations need to ensure that the third-party vendors they do business with are as equally protected against current and emerging cyber threats as they are. This is especially true because, as we've seen, a successful ransomware attack against a single MSP or Cloud-Based Service Provider can be debilitating to all their customers, as well as to their own company, as the attack can quickly shut down key systems which the customers depend upon to run their organization."

And of course a ransomware attack against an MSP can be fatal, putting an MSP out of business, which appears to have happened with PM Consultants, which is an Oregon-based IT consulting and IT support provider to dental practices. We talked about the dental practices that got hit. After PM Consultants was hit by ransomware in early July, they just gave up. They shut down their business later that month, saying that they were doing so in part due to a devastating event which had befallen their organization. And we know that their clients were all targets and victims of ransomware.

So our takeaway here is that - and this is the point for our listeners that I think is important. Any enterprise who is contracting for the services of any outside provider should give serious consideration to the sort of access that provider truly needs to have into their network. It's entirely natural for any provider to have a full trust in their own capabilities and their own security. And so they could in full good conscience ask for full and unfettered access into their clients' networks. And based on the infections we're seeing, that's apparently what they're being given. But any highly responsible IT manager who is on the receiving end of such an arrangement should give some serious consideration to exactly what sort of access the outside party needs, understanding that something entirely inadvertent on the other end could happen.

And, you know, it's not possible to offer any more concrete advice without knowing exactly what the relationship is and what form of access is really needed. But we're seeing this new pattern where managed service providers are targets because it allows the malware providers to essentially explode into their entire customer base and get a huge multiplicative benefit from a single infection. And so from my standpoint it's a little bit like the firewalls of yesteryear and today. Initially, when you would hook your LAN to the Internet, you would just hook your LAN to the Internet. It's like, yay, you know, we're a big friendly globe; right? Well, of course, not any longer. Now our firewalls are buttoned down tight. And in fact we've talked about this before. Original firewalls were default open, and then they would block specific ports they didn't want to give bad guys access to. There isn't a single firewall around now that uses that logic.

JASON: It's the opposite now.

Steve: Yes. They are locked down. They are default closed. And then you punch little holes, hopefully little pinpricks through the firewall, only to allow a trickle of traffic, in the best case, from specific other sites and IPs that you want to trust, rather than everybody. A website, of course, doesn't have that luxury. They typically need to be open to the world. But so if you have a relationship with a managed service provider, even allowing access from only their IP, that's not going to help you in this case because it's their IP that will be downloading malware onto your network. You really need to look at the kind of access they need so that they're just not given carte blanche remote control and full software download capability into your network.

Again, without knowing in detail what's going on at, like, the nature of the service being provided, it's impossible to be more clear. But really, I would urge enterprises that have relationships with third parties to look at what would happen if the third party were controlled by malicious agencies because that's happening. And they represent a high-value target. Bad guys are going to work hard to get a managed service provider infected. And you don't want your MSPs, any of your MSPs, to be hosts of ransomware. So just to add that caution to everyone's radar.

JASON: I love this next story, by the way.

Steve: Oh, my god, yes. It's such a great story. So five months after returning his rental car, the brief renter of the car still has full remote control. He can track the vehicle, remotely lock and unlock it, and start and stop its engine. So the story is that when Masamba Sinclair rented a Ford Expedition from Enterprise last May, he was excited to connect with its FordPass. The FordPass app allows drivers to use their phones to remotely start and stop the engine, lock and unlock the doors, and track the vehicle's precise location. What could possibly go wrong?

So he's 34 years old, and he told the people, I think it was Ars who was doing the reporting on this, he said: "I enjoyed it, and logged into FordPass to be able to access vehicle features from my phone such as locking and unlocking and starting the engine. I liked the idea of it more than I found it useful." He says: "The UI looks good and works well."

So today Sinclair's opinion of mobile apps and rental cars is decidedly less favorable. That's because, five months after he returned the vehicle on May 31st, his app continues to have control over the car. Despite multiple other people having since rented the SUV in the intervening months, FordPass still allows Sinclair to track the location of the vehicle, lock and unlock it, start and stop its engine. Sinclair has brought the matter to Ford's attention, both through its website and multiple times on Twitter. So far Ford has done nothing to end his access.

He said: "All it took for me to initially connect was to download the app and enter the VIN" - that's, you know, the VIN number on the engine - "then confirming connectivity through the infotainment system. There might be a way to disassociate my phone from the car itself," he says, "but that hasn't happened yet. And it's crazy to put the onus on renters to have to do that anyway." He says: "I have had no problems at all and have even unlocked the doors and started the engine when I could see that the vehicle was at the Missoula Airport rental parking lot." Wow.

So we have in the show notes three of his tweets, the first on June 4th. He says to @Ford, he says: "I can still track and unlock the Expedition that I rented last week via the FordPass app. Huge safety concern for all future renters. I submitted a solution via Ford New Ideas" - yeah, there's a new idea - "to solve this, and it was denied. THIS NEEDS TO BE FIXED," he has in all caps. And then he posts a pic.

The next day he tweets: "It's day five since I returned my rental, and now someone else has rented it out. Do I need to start remotely unlocking it until they also start to complain? Please fix this!"

And it looks like a week later, June 14th: "I returned this car" - he's tweeting to Ford. "I returned this car two weeks ago, and you've shown no willingness to allow rental companies to remove my access to unlock it and start the engine. Maybe I'll just start randomly unlocking it." Hopefully he doesn't.

JASON: Hopefully, yeah. Don't do that.

Steve: Yeah, don't do that. FordPass is offered, as we know, by the Ford Motor Company and is available for both iOS and Android devices. It's one of several apps for connecting to Ford vehicles. The less-than-intuitive means for unpairing a vehicle and phone, not to mention the difficulty in knowing a device remains connected, represents a serious security and privacy risk, not to mention to renters, but to people buying a vehicle secondhand. Imagine you buy it from CarMax or something, and don't know or particularly care that the previous owner still has access. While Ford said infotainment screens will indicate when a device is paired, it's obvious that multiple Enterprise employees and renters - okay, this was - remember this is, what, six months ago. Multiple Enterprise employees and renters have continued to miss the warning. Even now, after the reporter of this discussed the problem with both Enterprise and Ford representatives, Sinclair's access still has not been revoked.

Sinclair said: "I've been opening the app and tracking the vehicle almost every day to see if my access is still there; and, sure enough, I can see exactly where my old rental, affectionately named 'The Beast,'" he wrote, "is at any given moment. This means that I can not only find this rental car whenever I want, but I can also unlock the doors and help myself to anything that might be inside."

Since proximity - and this is me thinking now. Since proximity to the infotainment system was required initially to complete the initial pairing authorization and authentication, and since occasional proximity would be an expected characteristic for any actual car owner, it would seem to me that a useful security tradeoff would be to have a device be forgotten if it hasn't been within physical proximity of the vehicle's infotainment system for, I don't know, what, a week? Or two? I mean, it sounds like repairing isn't that burdensome.

And, you know, a car owner, they're going to be in their car every day. Or maybe they only drive it on the weekends. Okay, so make it two weeks. And after two weeks forget the thing. Or maybe after one week, like, flash a warning and say, you know, "This device has not been near the car for the last week. Please confirm you want to keep it paired." So make it an affirmative statement of pairing endurance. And if you don't do that, it unpairs. I mean, it's nuts.

JASON: Yeah, I'm curious about this FordPass thing because this actually reminds me of like a question I've had when I do rent a car, and they have the infotainment system, and I want to play my music through the Bluetooth. So that requires pairing. So then as a renter you end up pairing to this thing. And maybe you choose not to share your contacts with the system.

Steve: I hope.

JASON: I usually don't share the contacts because that seems like a really bad idea. But I'm sure a lot of people just say, yeah, whatever, connect it; and then they don't think to remove it. So, I mean, I guess there should be some sort of mode in these things for rental car agencies for them to - on their turnaround. Because even what you're talking about, Steve, is somebody rents it the next day, that still doesn't prevent this other third party from locking the doors, turning off the vehicle at the wrong time or whatever. It's a liability. There needs to be some easy way for them to be able to deactivate that, if that doesn't exist already. Maybe it does, and they're just not using it.

Steve: Well, or in the rental mode it would be reasonable to turn down the "forget the device" delay to one day.

JASON: Right.

Steve: Because, you know, anyone who's renting a car is in it, like, you know, a lot. So, and if not, oh, boohoo if your device - I mean, besides, who wants to unlock their doors remotely? I mean, okay, maybe. Who wants to start the engine remotely? Uh, okay.

JASON: Maybe if it's cold outside. I can understand that one.

Steve: These are all, like, gee whiz doohickey, you know, I have it because the...

JASON: Yeah, not necessity; right.

Steve: Because I'm on the Internet, you know, my car is an IoT. Well, we already know what a bad idea that is.

JASON: It's very, very interesting that it's taken this long and still nothing.

Steve: Wow, yeah.

JASON: Just seems like the wires are crossed on that one because I would imagine if the right person at Ford knew of that...

Steve: Oh, the wires are short-circuited, Jason, no question.

JASON: Yup, indeed.

Steve: So Chinese-made drones in the U.S. are being grounded. I have no way of independently assessing the danger that foreign-made drones could pose to U.S. domestic security. But they are flying everywhere, and they are essentially flying video cameras. And they're connected to the Internet. And they're sending data back to China.

So what's in the news is that, pending a security review of these Chinese-made drones, the U.S. Department of the Interior, the DOI, announced last Wednesday that it is grounding all Chinese-made drones and drones with Chinese-made parts until it reviews its drone program. And the Department of the Interior turns out to have quite a program. But because it's simply not practical, that decision does not apply to drones, quote, "currently being utilized for emergency purposes such as fighting wildfires, search and rescue, and dealing with natural disasters that may threaten life or property."

So in recent years the Department of the Interior has enthusiastically embraced drones, publicizing the wide variety of ways it has deployed them. In addition to being deployed for emergency rescues and disaster monitoring, which I think is a cool application, they're used in more expansive long-term projects such as geological surveys and wildlife population monitoring. According to a 2018 report about its use of drones, the department owned at the time 531 drones and conducted more than 10,000 flights across 42 states and territories. Okay. So 365 days in a year, that's three a day. So, I mean, they're in use.

Other more recent reports - wait, no, 30 a day. I dropped a digit. Yes, 10,000 flights on 365 days, that's 30 a day. So, yeah. Other more recent reports have updated that number to 800 operating drones. So since their report in 2018 they've purchased another, what, nearly 300. The report did not specify what percentage of those drones were Chinese-made. However, the most popular and capable drones we know are made by the Chinese company DJI. And in fact two years ago the U.S. Army discontinued the use of drones produced by DJI, who is the world's biggest manufacturer of drones, because of the risk of these vulnerabilities.

And last May the DHS, the U.S. Department of Homeland Security, warned companies that Chinese-made drones could potentially transmit sensitive footage or data to third

parties. And then last month a bipartisan group of lawmakers introduced legislation that would bar all federal agencies from operating drones manufactured or assembled in China. Because the Interior Department uses drones to survey a variety of critical infrastructure, including mines and dams, as well as to study rapid response situations and emergency routes - and of course we know that these things have GPS in them also so they know what their location is - the information they collect has at least some potential for abuse.

So far to date, scant public evidence exists that Chinese drones have been involved in any large-scale cyberespionage, or that they have backdoors built in that would allow them to be exploited for surveillance. For their part, DJI has pushed back and argued that opposition to its products is motivated primarily by political hostility toward Chinese companies. But we all know that, just as with itty-bitty malicious chips hidden on motherboards, even if it isn't being done, it could be done.

So I suppose I'm glad that at least a sober awareness of the risks exist and that they're being taken seriously. DJI has proposed assembling more of its drones in the United States as well as offering a version of its devices called the "Government Edition," with certain safeguards that would protect information captured and would prevent it from being transmitted wirelessly. So, you know, just brainstorming a bit, what may develop would be a two-tier system where governments which can afford to pay for more costly, audited, and known clean drones will have that option.

But regular users like wedding photographers can save their money since recording Jack and Jill's nuptials from a height is unlikely to be of great interest to foreign spies. And DJI could continue to make their lovely hardware, but allow an open source community to add the software, if that is even being done. I know that there are a lot of drones that have been reverse engineered from a hardware standpoint and where there is now third-party open source software. So, you know, it just means let's get responsible here as government purchasing agents. If DJI wanted to salvage their reputation, Jason, they could just make it official, allow the government to take their lovely hardware and put their own firmware in it so they know exactly what it's doing, not make it a closed black box, or in this case a cream-colored four-armed shape.

JASON: With deadly propellers on top. Yes, indeed, indeed.

Steve: Yeah. They are nice drones.

JASON: All right. Let's get into the controversy here, Steve.

Steve: So as I mentioned at the top of the show, DoH remains controversial, not so much with end users, who appreciate the idea that their ISPs and others who might wish to snoop on and possibly filter and block their DNS lookups, will be much less able to do so. But within the industry, where those who are being blinded to DNS lookups are crying foul and running around stirring up a bunch of FUD - you know, Fear, Uncertainty, and Doubt - it's sort of breathtaking, like, how much kerfuffle there is. For example, last week Motherboard carried a story with a headline: "Comcast Is Lobbying Against Encryption That Could Prevent It From Learning Your Browsing History."

JASON: I remember that, yeah.

Steve: So, yeah, so that's what Motherboard is saying. Comcast is lobbying against encryption, that is to say, wants laws to prevent DNS encryption that would prevent it from obtaining its browser's history. The subhead reads: "Motherboard has obtained a leaked presentation which Internet service providers are using to try to lobby lawmakers against a form of encrypted browsing data." They put that in there to simplify things.

They said, okay, so anyway, for our audience let's remember, as we've previously covered here, Mozilla has indicated that it plans to have Firefox route its DoH queries to Cloudflare's servers.

In response to the flack it's been getting, Mozilla has been explicit that "no money is being exchanged to route DNS requests to Cloudflare" as part of their DNS-over-HTTPS feature that's currently being gradually enabled for Firefox users in the U.S. And, for example, they further indicated that the only reason they chose Cloudflare is that they are being rigorous about the requirements, the privacy requirements and the no-profit-from-this-data requirements, of anyone that they point their Firefox data to. And as other DoH providers become willing to make the same pledge that Cloudflare has, then Firefox is happy to use those, too. I mean, this is so, you know, it's important to understand we're just at the beginning of this.

So detractors have been saying that by using Cloudflare as the default DoH resolver for Firefox, Mozilla will help centralize a large chunk of DNS traffic on Cloudflare's service. And that is indeed the case. Let's also remember, as for Google, Google is planning to take a much more dynamic posture. We've covered this previously. Google will have Chrome test the user's currently configured DNS server. And if that service also supports DoH, Chrome will switch over to using it. This is, I think, a slick solution, since Chrome users may have already manually switched their non-encrypted DNS to some other provider who offers superior value-added services for DNS that they want. So if that's the case, why not allow it to also be encrypted in addition to having those value-added services.

So, okay. With that in mind, and Comcast knows all of this, let's look at what Comcast is alleging to lawmakers. Motherboard wrote: "The plan, which Google intends to implement soon, would enforce the encryption of DNS data made using Chrome, meaning the sites you visit. Privacy advocates," writes Motherboard, "have praised Google's move. But ISPs are pushing back as part of a wider lobbying effort against encrypted DNS, according to the presentation" which Motherboard obtained. "Technologists and activists say this encryption would make it harder for ISPs to leverage data for things such as targeted advertising, as well as block some forms of censorship used by authoritarian regimes."

Mozilla, they write, which makes Firefox, is also planning a version of this encryption. Marshall Erwin, senior director of trust and safety at Mozilla, told Motherboard in a phone call after reviewing sections of the Comcast slide deck that "the slides overall are extremely misleading and inaccurate. And frankly," he said, "I would be somewhat embarrassed if my team had provided that slide deck to policymakers." And he added, "We're trying to essentially shift the power to collect and monetize people's data away from ISPs and providing users with control and a set of default protections."

So they had one summary slide that I'll share. And I'm just going to just read the headlines of the other slides. I think there were 17 total. So their summary that begins this says: "Google has announced unilateral plans" - that's in all bold - "unilateral plans, along with Mozilla, which derives over 90% of its revenue from Google." Now, I don't think that's true anymore. Remember that that used to be the case, that Mozilla was - I think they defaulted to Google Search. I mean, we know that historically Mozilla had a relationship with Google. But I think I remember Leo telling me that that ended years ago. So that seems, if that's correct, then wow, crazy to put that in there. Anyway, unilateral plans to activate DoH in its Chrome browser as soon as October, that is, you know, now.

"Google also appears poised to activate DoT for devices using its Android mobile operating system in the near future." And of course we covered that, as well, at the time. "If activated, this feature would by default route all DNS traffic from Chrome and Android

users to Google Public DNS." Of course that's not true. Thus, oh, we've got a big bold chunk here, "centralizing a majority of worldwide DNS data with Google."

And then they said "This charge would mark a" - here we are in bold - "fundamental shift in the decentralized nature of the Internet's architecture and give one provider control of Internet traffic routing and vast amounts of new data about" - one provider control of Internet traffic routing. What? "And vast amounts of new data about consumers and competitors." Uh-huh. Okay. Except that Chrome is like, you know, there's like a list of nine non-Google Public DNS servers that they work with out of the box.

JASON: Don't say that. No one needs to know that.

Steve: Yeah.

JASON: There's nothing to see here.

Steve: "The unilateral centralization," yeah, "of DNS raises serious policy issues relating to cybersecurity." That's in bold. "Privacy, antitrust" - oh, those are not bold. Oh, "national security and law enforcement." Those are in bold. "Network performance and service quality, including 5G and other areas." So, okay. In the presentation Comcast paints this type of encryption as something that will fundamentally change the Internet and will centralize power under Google.

Comcast wrote in its presentation: "The unilateral" - oh, this was in the slide that I just read. Okay. So I have a link to the PDF of the slide deck for anyone who really wants more pain. But here are the 11 titles of slides. The presentation itself is titled "Google's Proposed Public DNS Plans." Slide one, or one of 11: "Google's unilateral imposition of default centralized DNS encryption will harm key components of the Internet." Okay, slide two. "Google's plans will cause radical disruption." Oh, geez.

Number three: "Significant cybersecurity and national security risks." Okay, what? Four: "Privacy risks increase as a single firm amasses more consumer data." Wait. Okay. More than Comcast getting their hands on it? Okay. Let's remember that back in 2017 ISPs including Comcast aggressively lobbied Congress to make it possible to sell their users', their customers' browsing data without their consent or knowledge. Okay. Five: "Antitrust and competition concerns." Wait. How is DNS a competitive territory? It maps domain names to IPs. Six: "Law enforcement efforts may be compromised." Okay.

Seven: "CDN localization will likely suffer, and backbone costs will rise." Okay, now, I would give them that one. CDNs do rely on localization, that is, it can be the case that the IP you obtain varies. But more often it's the CDN localization is performed at the router level, and the IP you get never changes. So, sorry, number seven. Number eight: "Wireless performance, including 5G, will be undermined." No, it won't. Nine: "Parental controls and content filtering may be disabled." Okay, yes. If the DNS provider you used was using DNS for parental control or content filtering, and you switched to a different DNS provider that doesn't offer that, then yeah, that will change. But if your DNS provider offers DoH, Chrome will detect that automatically and just encrypt your queries so that your ISP can't see them, but your service is not otherwise affected. Ten: "ISPs and other enterprise services may be disrupted or broken." What?

JASON: How?

Steve: ISPs and other enterprise services may be, well, okay, there is the issue of enterprises. But again, Chrome has already anticipated this and will not override non-public DNS queries, so enterprise services continue to be functional. And 11, the last one: "Congress should demand the Google pause and answer key questions."

So we've really got to ask ourselves why Comcast apparently places so much value upon their access to their customers' DNS data. You know, they've got their non-DNS data. They've got all their data. So wow. They really do seem to have their panties in a bunch over this thing.

JASON: That data's as valuable to them in their eyes as it is to Google or anyone else, I suppose.

Steve: Yeah.

JASON: Yeah, now would be the time to strike as far as that's concerned.

Steve: Well, they're doing that, yeah.

JASON: Yeah, in many ways.

Steve: So, and speaking of - nice segue, Jason - time to strike, it took a long time, but the BlueKeep-based attacks have finally started. And what we predicted on this podcast has finally happened. The so-called BlueKeep vulnerability in Windows Remote Desktop is being exploited and, as we suspected, not by a worm. ZDNet's headline was "BlueKeep attacks are happening, but it's not a worm. Hackers are using BlueKeep to break into Windows systems and install a cryptocurrency miner." Also what we expected.

The Hacker News says: "First Cyber Attack 'Mass Exploiting' BlueKeep RDP Flaw Spotted in the Wild." Threatpost: "BlueKeep Attacks Have Arrived, Are Initially Underwhelming." And BleepingComputer: "Windows BlueKeep RDP Attacks Are Here, Infecting with Miners." So some friends of ours are in the middle of this. Over the weekend and across the world, security researchers' honeypots monitoring port 3389 - which is the RDP port, 3389, which is where Windows listens for incoming remote desktop protocol connections - the honeypots began lighting up, actually, and crashing as attempts were being made to leverage the BlueKeep vulnerability for the purpose of running cryptocurrency miners.

On Saturday Kevin Beaumont noticed that multiple honeypots in his EternalPot RDP honeypot network started to crash and reboot. His pots have been active for almost a year, so they would have been scanned and catalogued by anyone who was planning an attack. And during all that time, this is the first time they came down.

The first details about BlueKeep being the cause of these events came from our friend Marcus Hutchins (a.k.a. MalwareTech), who investigated the crash dumps from Kevin's machines. He said that he "found BlueKeep artifacts in memory and shellcode to drop a Monero miner." According to Marcus's analysis, an initial payload runs an encoded PowerShell command that downloads a second PowerShell script, which is also encoded. Marcus says that the final payload is a cryptocurrency miner, likely Monero, currently detected by 25 out of 68 AV engines at VirusTotal.

Marcus said that the malware may not be a worm, but that it is mass-exploiting the BlueKeep bug. This indicates that the cybercriminals are using a BlueKeep scanner to find vulnerable systems exposed on the 'Net and drop the cryptocurrency miner on them. And in a later update Marcus said that analysis of its network traffic does not indicate self-propagation.

So the server doing the exploitation obtained its target IP addresses before the attack from a predefined list. This is what we expected, since the exploit is straightforward once its details have been worked out. And adding worm code just adds additional unnecessary complexity without returning any real value. Once upon a time, when the Internet was a much more quiet place, a worm would have been launched as a means of obscuring the original attacker's origin. But on today's Internet, where other machines

are readily compromised and used to launch reflected attacks, hiding is far less necessary. The first public BlueKeep exploit was added to Metasploit two months ago in September. And Marcus's analysis has confirmed that the same code which was present in the open source Metasploit module is also present in the malware.

In other words, it looks very much as though whoever it was who was behind the attacks did not independently develop the attack, but is using publicly available code resources and made no effort to turn it into a wormable threat, as Kevin's honeypot crashes suggest. It's not particularly reliable, either, or it wouldn't be crashing these things, it would be installing itself. Kevin recently noted that there are presently more than 724,000 systems worldwide that are currently susceptible to BlueKeep exploitation. Three quarters of a million machines now on the 'Net, you know, this thing got closed by Windows Update months ago. These systems just never update themselves. They've got RDP exposed and just waiting to get taken over.

So anyway, BlueKeep has happened. It's going to be Monero. Someone's going to be making money from Monero mining for them on all these machines that presumably no one is paying any attention to because they're sure not updating them with any useful code updates.

JASON: There's just too much to pay attention to. Your show is always just so full of joy and jolly, jolly insight into just how safe the Internet really is. We're always talking about millions and millions of vulnerable systems. It's very cheerful.

Steve: Get your tinfoil.

JASON: Steve, excellent stuff here. And always bringing so much knowledge on these topics. Really appreciate the work that you do, and keeping people safe, helping keep people safe, anyways. You mentioned it earlier. We should mention it again. The SQRL event is coming up Saturday, November 30th. And now just a few weeks away because time is weird and going very fast. So November 30th, 2:15 p.m. Pacific.

I'm not sure if there are any openings still for this. I think last I checked it was pretty full, but there were a couple of open spaces. So anybody who wants to sit in the studio for that recording is welcome to email tickets@twit.tv. And only do that if you plan, if you're like 100% I will be there. I will be sitting in the audience. That way someone else who could make it isn't, you know, that opportunity isn't taken from them. So we definitely want to make sure that people who sign up to be here can actually make it. So if you can, tickets@twit.tv. That's Saturday, November 30th, 2:15 p.m. Pacific for the SQRL event.

Steve: And I would just add that, if your plans change, and you cannot make it, but you did promise to be there, please also do the same courtesy and just let TWiT know that, whoops, you know, you really wish you could make it, but just, you know, can't get away.

JASON: Yeah. It's a great, great point. Following up would be awesome in that regard. GRC.com is where you can go for everything Steve, everything that Steve talks about on this show and offers. Of course SpinRite, which you know all about. SQRL, like we were talking about, if you want information about SQRL, and you're like, hey, now that I've read up on it, now that I know a little more, I'll go to the event. You can find that information at GRC.com. Also audio and video of Security Now! can be found there, as well as the transcripts of this and other episodes can be found there, as well. That's the only place that you can find the transcripts. So if that's what you're looking for, GRC.com.

Otherwise, what you're seeing right now is our website, TWiT.tv/sn. This is the Security Now! show page at TWiT.tv. And here you can also find all the audio and video of every episode of Security Now!; ways to subscribe, which is really what you should do. Just go there, click on the link to your podcatcher of choice, and you'll be subscribed. You won't even have to seek out the episodes. They get delivered to you through podcast magic.

We record this show every Tuesday, live, 1:30 p.m. Pacific, 4:30 p.m. Eastern, and we've had a time change this last Sunday. So is it 20:30 UTC or 19:30 UTC? I don't actually know. I forgot to look into this before starting the show. But check your clock is the best that I can do on short notice.

Thank you, Steve. Always fun doing the show with you and appreciate the opportunity.

Steve: I appreciate it, and I had a great time, and we'll do it again next week, my friend.

JASON: Absolutely. Sounds good. And also thanks to Jeff behind the board. We'll see you next week on Security Now!. Bye, everybody.

Steve: Bye.