Good Samaritans?

Description: This week we observe the untimely death of Microsoft's co-founder Paul Allen, revisit the controversial Bloomberg China supply chain hacking report, catch up on Microsoft's October patching fiasco, follow up on Facebook's privacy breach, look at the end of TLS v1.0 and 1.1, explore Google's addition of control flow integrity to Android 9, look at a GAO report about the state of U.S. DOD weapons cybersecurity, consider the EOL of PHP 5.x chain, take a quick look at an AV comparison test, entertain a few bits of feedback from our listeners, and then consider the implications of grey hat vigilante hacking of others' routers.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-685.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-685-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with lots of security news. DuckDuckGo is on the rise. Details about the Microsoft patch that deleted data. Yikes. And which antivirus is the best? Steve has some information. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 685, recorded Tuesday, October 16th, 2018: Good Samaritans?

It's time for Security Now!, the show where we cover your privacy and security online. We also talk about all the latest security news, some tech news, maybe even explain how some of this stuff works, thanks to this guy right here, the Explainer in Chief, Steve Gibson. Hello, Steve.

Steve Gibson: Mr. Laporte, it's good to be back with you again this week for our 685th episode.

Leo: Wow.

Steve: Yeah. So the thing that caught my interest, and I want to do this more as a discussion, just sort of of the pros and cons, because I think it will be interesting. My promotion of the idea of auto-updating routers has generated some controversy. But this particular news will probably generate additional controversy. So controversy is good. That's why it's the title of the show: "Good Samaritans?" with a question mark. It turns out somewhere, we don't know where - oh, wait, I think he's Russian speaking.

Leo: Oh, well, there you go.

Steve: There's a Russian-speaking Good Samaritan, question mark, who claims to have, and there's been some confirmation, so far patched more than 100,000 of the vulnerable MikroTik routers without their owners' knowledge or permission, just saying, well, I'm going to fix this for you.

Leo: Taking advantage, I presume, of the exploit in the first place, right, to get in there to fix it.

Steve: Yes, yes, exactly. And then essentially what he's doing, he's not updating firmware, he's installing firewall rules to prevent additional exploitation. And there's a lot of sides to this because, well, and we'll discuss it without...

Leo: Interesting.

Steve: ...stepping on our own lede. We're also going to talk about, just briefly note, as you did at the end of MacBreak Weekly, the death, sad death - I don't know if death is ever good. In this case it's sad. Microsoft's cofounder Paul Allen passed away yesterday at the age of 65. Also we're going to revisit some of the things that we've been talking about recently, just for some follow-ups. The Bloomberg China supply chain hacking report, which is continuing to be controversial. There is some additional news on Microsoft's October patching fiasco. It's hard to characterize it as anything but.

Facebook revised downwards the number of people who were breached and had their private stuff stolen from them. And oh, my god, Lawrence Abrams over at Bleeping Computer said something that was so funny. Every time I think about their revision it makes me chuckle, so we'll share that. We're going to look at the end of TLS versions 1.0 and 1.1; explore Google's addition of something known as "control flow integrity," which was added for the first time to Android 9; take a look at a GAO report about the very sad state of U.S. Department of Defense weapons cybersecurity that actually had some bit of humor in it. Well, yeah, it did. I was trying to think whether that came from reporting or from the report itself; but, no, it was in the PDF.

We're going to look at the implications of the forthcoming end of life of PHP 5.x. That's at the end of this year, and that could be a problem. We'll take a quick look at an AV comparison test, recognizing as we always do that they're always a little suspect because you've got to wonder, well, who contracted for that? But it was interesting. We're going to, if we have time, entertain a few bits of feedback from our listeners, including one from John McAfee, who I don't think is a listener because he would be smarter if he was a listener.

And then we're going to consider, as I mentioned, the implications of grey hat vigilante hacking of other people's routers, ostensibly or by intention to fix them. And then of course we have our typical fun Picture of the Week. So I think another great podcast to entertain and engage our smart listeners who've been following the podcast for years.

Leo: Jammed with great stuff. Steve?

Steve: So our Picture of the Week raises a few questions. First of all, foremost among those questions is, is this progress? Does this represent progress? We're looking at a Coca-Cola machine where the screen is explaining that "This Dispenser is temporarily Out of Operation. Please wait while the dispenser performs required nightly maintenance." And then in all caps, which makes me wonder who this message is aimed at: "DO NOT SHUT DOWN THE DISPENSER."

Leo: Don't unplug it, whatever you do.

Steve: And then there's a sort of a clock-y kind of progress thing. The guy who sent this to me, I responded via Twitter, and I said, "Hey, thanks. It'll be a great photo for tomorrow's podcast," because this was yesterday. And he wrote back, and he said, "Just FYI, the progress meter was not moving." Now, the other worry is that it says "Downloading Updates." So it's like, okay. So your Coca-Cola dispenser has to have an Internet connection. Now, the screen is also glossy, so we can see the reflection, we sort of see the shadowy outline of the guy who took the picture. But behind him is Windows and daylight, which really makes this look like it's not night. And there's something that says "01:26:37." And I'm thinking, well, I hope that's not the remaining time.

Leo: I think it is.

Steve: Counting down.

Leo: In an hour and a half you'll get your Coke.

Steve: You can have your drink. And really, "Do not shut down the dispenser." So is this for the repairman? Or, I mean, really, a mischief-y scoundrel would be tempted to do just that because, you know, there's a plug back there somewhere. So I don't know. This whole thing, I'm skeptical about this kind of progress. I mean, for example, why is it that it can't be happily doling out Coca-Cola while downloading those updates in the background? Does it have to stop? It can't run its Coke dispenser while it's - there's no multitasking here or what? It just doesn't make any sense to me. So it's like, okay, fine. Were this my Coca-Cola machine, it would be downloading in the background. And then it would say, hold on one second while I reboot.

Leo: I'm guessing it's one of these newfangled Coke Freestyle machines where they take all the different syrups - see, it's got a big display on the front there. I bet you that's what this is.

Steve: Yeah, actually he did say that you could mix your Cola to your taste.

Leo: Yeah. It's a freestyle. So I guess you need an operating system to mix your Cola to taste? I don't know. It's not your average fountain unit.

Steve: Probably better than valves, Leo, because I think those would be...

Leo: Your valves are being updated.

Steve: Yes.

Leo: Yeah. I was thinking of getting one of those for the office, but maybe not now.

Steve: Well, I would love to have some feedback about what it's doing in the middle of the afternoon updating for an hour and a half when you've got thirsty coworkers sitting around here saying, "Hey, boy, Leo, this was a real improvement."

Leo: I think it's downloading the new Fanta Halloween flavors. They're scary good. That must be it.

Steve: Oh, my goodness. Okay, no, I'm not going to get into this any further.

Leo: You don't want this stuff. Trust me. You don't want this stuff.

Steve: I'm happy with my single-flavor espresso, thank you very much.

Leo: This machine does have an app, so I'm thinking, app-enabled, it probably needs updates from time to time.

Steve: Yeah. Again, it should just do it by itself.

Leo: In the middle of the night.

Steve: Well, or it should be able to give you foam and fizz and things while it's downloading to itself and then say, oh, hold on one second. It kind of makes a burp sound and then reboots. But no.

Leo: You can actually use the app to make your own mixes.

Steve: I'm sure you can.

Leo: And there's some way that you can go to the machine and share it with the machine.

Steve: What could possibly go wrong? So okay. Lots of controversy raised by Bloomberg's report - I notice I spelled their name wrong in my title, but I know that's an E-R-G - about the Chinese supply chain problem that was the topic of last week's podcast. And I was just curious to see what had happened since and what was the upshot and so forth. Business Insider reported that "All parties involved have denied the

report, including most recently the Secretary of the Department of Homeland Security during a Senate hearing [last Wednesday]."

She said, that's Kirstjen Nielsen said: "With respect to the article, we at DHS do not have any evidence that supports the article. We have no reason to doubt what the companies have said." And of course that's everybody involved denying it. But shortly after she said that, after saying there's no evidence, she said: "We can tell you, though, it's a very real and emerging threat that we're worried about." So we're going to be worried about it, but as far as we know - and I should also mention that Bloomberg, as one would expect, I guess they could retract if they believed that their story had fallen through. They're sticking by it.

Leo: Oh, no. They're not going to retract it. They worked on this for two years.

Steve: Yeah.

Leo: I talked to Mark Milian, who works at Bloomberg Business Week, on Sunday. He was on TWiT. Their editor-in-chief, who's highly respected, was at The Economist for a long time, was very much involved in vetting the story. I mean, this was a cover story for the Bloomberg Business Week magazine. They're not going with this unless they have absolutely every confidence it's true.

Steve: Yes, yes. And what I got a kick out of was that during that same hearing Christopher Wray, who's our current FBI director, said something I thought was curious. He said he couldn't confirm nor deny the existence of any investigation into compromised Supermicro equipment. And I'm thinking, why couldn't he simply deny it? You know? I can neither confirm nor deny. What is that?

Anyway, so TechCrunch's Zack Whittaker, who as we know has been reporting on security for years, considers this confusion of press releases and statements to be the nature of this kind of reporting. When the stakes are this high, I mean, and they are, reporting on security vulnerabilities and classified information means that you're more often than not dealing with making your sources anonymous to protect them, which rightfully opens, I mean, and of necessity opens your work up to denials and condemnations from the organizations you cover.

And we have to remember also that here in the U.S., companies who are legally forced to divulge information to law enforcement are also bound to deny that any such disclosure ever took place, you know, the famous FISA warrant mess. So whatever happened occurred several months ago, as we know, dating at least from back August. You indicated that they've been working on this for quite a while.

Leo: This goes back to the Obama administration. This has been a long investigation.

Steve: Yeah, yeah. So while, yes, we don't have any absolute proof, with something that is this high stakes, we're not going to. I mean, there isn't going to be a smoking gun. Everybody had to put their guns away. They've hidden them all. And they said, what gun? No, there's no gun here. So anyway, I just wanted to follow up that, yes, it generated of necessity some reaction from the world. But anyway, I'm glad you had him on, and you guys discussed it further.

Leo: Yeah. And I trust Mark. He knows these guys. He shared a desk with Jordan. He says no, no, there's no question that these are good sources and that this is a good story.

Steve: And I also saw, sort of apropos of this, Kevin Mitnick just demonstrated a malicious USB cable which was entering keystrokes into a fully patched Windows 10 machine to install malware.

Leo: The NSA had such a thing, according to Snowden.

Steve: Yes, yes. And when President Trump and his entourage went to meet with the leader of North Korea, remember they were handing out little USB fans, and everybody was quite worried about, wait a minute, how hot am I? Do I really need to...

Leo: How hot is this fan?

Steve: Yes. Do I really need to cool myself off that much? Okay. So also following up on our coverage of Facebook from last week, as we know, they originally acknowledged a sophisticated attack which leveraged a vulnerability in a video uploading tool and their "View As" feature, which allowed essentially a pivot attack where an attacker could obtain the authentication for any other targeted user. So they have posted an update on the 12th of October titled "An Update on the Security Issue." I've got a link in the show notes for anyone who wants it in every detail.

But what cracked me up was that Lawrence Abrams covered this for his Bleeping Computer site. And so quoting from Facebook's update, he first said: "We now know that fewer people were impacted" - this is Facebook saying, "We now know that fewer people were impacted than we originally thought. Of the 50 million people whose access tokens we believed were affected, about 30 million actually had their tokens stolen." To which Lawrence wrote: "Isn't that great? Only 30 million."

Leo: Only 30 million actually stolen. That's what's amazing.

Steve: Yes.

Leo: That's mindboggling.

Steve: Yes. So Facebook wrote, so that they did provide additional information: "First, the attackers already controlled a set of accounts." Which, you know, who doesn't have a Facebook page? I mean, even I have one, just to go check on security features. "First, the attackers already controlled a set of accounts, which were connected to Facebook friends. They used an automated technique to move from account to account so they could steal the access tokens of those friends" - that's what I call "pivoting" from account to account - "and for friends of those friends, and so on" - so like a big network tree that branches out - "totaling about 400,000 people.

"In the process, however" - this is Facebook writing. "In the process, however, this technique automatically loaded those accounts' Facebook profiles, mirroring what those 400,000 people would have seen when looking at their own profiles. That includes posts on their timelines, their lists of friends, groups they are members of, and the names of recent Messenger conversations. Message content was not available to the attackers, with one exception. If a person in this group was a page admin whose page had received a message from someone on Facebook, the content of that message was available to the attackers."

Second paragraph: "The attackers used a portion of these 400,000 people's lists of friends to steal access tokens for about 30 million people." That is another stage of pivot out from all of the 400,000. So this thing just sort of escalates as it branches out. For 15 million of those 30 million, attackers accessed two sets of information: name and contact details, so phone number, email, or both, depending on what people had on their profiles.

For the other 14 million, the attackers accessed the same two sets of information, as well as other details people had on their profiles. This included username, gender, locale, language, relationship status, religion, hometown, self-reported current city, birth date, device types used to access Facebook, education, work, the last 10 places they checked into or were tagged in, website, people or pages they follow, and the 15 most recent searches. So a lot of information on 14 million. And then they finally said that for that remaining one million people the attackers did not access any information.

So I will reiterate something that all of this information has made clear and I mentioned last week, which is - but I haven't seen it noted anywhere else. And that is that this was also a fully targetable attack. I worry that the size of the number sort of leaves people to believe that this attack was just sort of like a "you get what you get" sort of attack. But first of all, we know that both of the top two Facebook executives were compromised. Maybe they just fell into this 30 million, but also probably - because the idea is, since you could say "view as somebody else would see me," and then in the process you acquire that person's authentication token, you can ask for anybody you want to see what your page looks like and then pivot to their identity.

So to my way of thinking, while yes, 30 million is a big number, albeit down from 50, the fact that these attackers could get this information, although again not total account takeover, just viewable information, but also logged on authentication for anybody they wanted, that makes it in my mind more serious. And also missing from this Facebook update is they neglected to mention that the then owner of someone's authenticated Facebook identity could then use the widespread Sign-in with Facebook, OAuth, to log into many other Internet websites and services under the identity of that impersonated account, which again is very powerful and potent. But Facebook decided not to update us on that issue. So anyway, I just wanted to follow up when I read what Lawrence wrote: "Isn't that great? Only 30 million."

Leo: Oh, what a relief.

Steve: Yes, we could all, well, as we know. Oh, and they concluded their posting, Facebook did, by saying people can check whether they were affected by visiting our Help Center. In the coming days we'll send customized messages to the 30 million people affected to explain what information the attackers might have accessed, as well as steps they can take to help protect themselves, including from suspicious emails, text messages, or calls. And it's not a simple URL: Facebook.com/help/securitynotice, all one word, ?ref=sec.

Leo: They wouldn't want you to find it too easily.

Steve: But that will take you to the page.

Leo: You could probably go to their security page and navigate from there, I guess.

Steve: Yeah, and dig in, yeah. Okay. So also an update on the Windows 10 October update mess. And actually there was so much that has happened in a week, it has been such a debacle that I found it a bit challenging even to organize it into a coherent view for this podcast. So last Tuesday afternoon, while we were recording last week's podcast, Microsoft posted to the Windows blog "Updated version of Windows 10 October 2018 Update released to Windows Insiders." Okay? Not the general public yet.

So they said: "Last week we paused the rollout of the Windows 10 October 2018 update version 1809" - so that's the feature update - "for all users as we investigated isolated reports of users missing files after updating. Given the serious nature of any data loss" - okay, now, get a load of this compared to how they end with and something we learned subsequently. "Given the serious nature of any data loss, we took the added precaution of pulling all 1809 media across all channels, including Windows Server 2019 and IoT equivalents. We intentionally start each feature update rollout slowly [huh?], closely monitoring feedback [huh?], before offering the update more broadly." And of course in this case they're really glad.

"In this case," they said, "the update was only available to those who manually clicked on 'check for updates' in Windows settings. At just two days into the rollout, when we paused, the number of customers taking the October 2018 update was limited. While the reports of actual data loss are few" - and here they quote the number you mentioned last week, Leo - "one one-hundredth of 1% of version 1809 installs, any data loss is serious."

Okay, so yes. And what Microsoft failed to note here was the sobering news that there had previously been many reports from Microsoft's own Win10 insiders of exactly this mass data deletion occurring to them which Microsoft had ignored. So that slipped through the cracks. Now, to Microsoft's credit - as we know, anybody can make a mistake - they are working to fix this.

They said: "Prior to re-releasing the October 2018 update, our engineering investigation determined that a very small number of users lost files during the October 2018 update. This occurred if Known Folder Redirection had been previously enabled." And Leo, you brought this up while we were covering this last week because that news had just occurred. "But files remain in the original 'old' folder location versus being moved to the new redirected location. KFR is the process of redirecting the known folders of Windows including Desktop, Documents, Pictures, Screenshots, Videos, Camera Roll, et cetera, from the default folder location" - which is c:\users\username, then folder name - "to a new folder location. In the previous feedback from the Windows 10 April 2018 update" - okay, so that's April 2018 update, which was the previous feature release.

"In the previous feedback from the Windows 10 April 2018 update, users with KFR" - that is, the Known Folder Redirection - "reported an extra empty copy of known folders on their device. Based on feedback from users, we introduced code in the October 2018 Update to remove these empty, duplicate known folders. That change, combined with another change to the update construction sequence, resulted in the deletion of the original old folder locations and their content, leaving only the new active folder intact."

So now we know in more detail exactly what happened. It was that in some cases the contents of the original known folders was not moved to the new relocated known folders, leaving that behind. So at that point, for the last six months, since the April update, users essentially had sort of abandoned orphaned documents in their original folder locations, and then the post-relocation new locations.

And then unfortunately, as a consequence of, you know, Microsoft isn't telling us exactly in what detail, but they didn't apparently look to - they didn't move the stranded orphaned documents then. They just killed the folders. Which caused a loss of data for the users who still had and were using the non-relocated known folder contents, which then disappeared. And they go on to explain how they found and fixed three scenarios where this was seen to occur. I've got it in the show notes for anyone who's interested. I won't drag us through it right now. Then they indirectly addressed the fact that insiders had been screaming about this problem well before the update's release and being ignored.

Microsoft concluded their posting saying: "To help us better detect issues like this, today we have enabled a new feature in the Windows Insider Feedback Hub. We've added an ability for users to also provide an indication of impact and severity when filing User Initiated Feedback. We expect this will allow us to better monitor the most impactful issues, even when feedback volume is low." And again, feeling a little chastised, I'm sure, they said: "We will continue to closely monitor the update and all related feedback and diagnostic data from our Windows Insider community with the utmost vigilance. Once we have confirmation that there is no further impact, we will move towards an official re-release of the Windows 10 October 2018 update. We apologize for any impact these issues may have had," blah blah blah.

So there were reports of it on the feedback hub as a consequence of, well, aided by the fact that there wasn't the ability before now for people posting feedback to really raise an alarm algorithmically. I mean, apparently lots of people were saying, "Oh, my god, files have just been deleted." But that slipped through the cracks. So now Microsoft has added the ability to specify the severity of the problem. So that's good. And hopefully that will help to prevent this moving forward.

Also since this October round of patches, in both the 1803, which is the current pre-feature update, and the 1809, which is the October update, both of those updates, a bad HP keyboard driver had somehow made its way into both 1803 and 1809 updates, which was causing a blue screen of death on boot for HP machines and was triggering a WDF, Windows Device Framework or Driver Framework violation. And this is of course problematical for users who are not technical because, if you installed the updates and then rebooted, at that point you were never able to get Windows up again. It would blue screen while it was trying to boot.

So it took intervention in order to back out of that driver. Microsoft explained that, if you hadn't rebooted, you could get rid of the driver, which was v11.0.3.1, and then you'd be okay for rebooting. But if you had rebooted, you need to go to additional measures. And there's Knowledge Base articles and so forth. I've got the link in the show notes for anyone who's interested. So Microsoft acknowledged that this was a problem that crept into this month's updates for Windows 10.

Also there is a "what needs your attention" screen which could come up and which might have been a blessing, actually, complaining that an Intel audio display notification problem existed, which was some sort of compatibility issue with a range of Intel display audio device drivers that Microsoft said might result in excessive processor demand and reduced battery life. As a result, the update process to Windows 10 October 2018, that is, the feature update to 1809, will fail.

If you see a "what needs your attention" notification when you run the October update, you have an Intel display audio device driver, and then they note which one it is, installed on your system that is preventing that update from occurring. Which, as I said, may have been good because, also and separately, people who did get the install, actually either the regular update, just regular monthly 1803 update, or the feature update, some people were reporting, independent of what their hardware was, that their sounds died. They were getting an Intel SST Audio Controller problem which was killing audio for those users. In that case, since this doesn't prevent your system from booting, you can go into Device Manager and delete this - it's called the Intel Smart Sound Technology, driver version 9.21.00.3755 - and that gets your sound back.

Also there was a display brightness resetting problem which is affecting some users such that every time they booted Windows 10 after the October updates, their brightness was set to minimum. No word on the resolution for that yet. And with the release of Windows 10 1809, which is the feature update, the Microsoft Edge web browser and Microsoft's UWP Store apps might no longer be able to connect to the Internet, which is a pesky problem for a web browser. It turns out that Edge and the UWP apps now require, as of 1809, TCP/IPv6 to be enabled, or they will not connect. And so I'm sure they're enabled by default. So presumably, if someone had previously manually disabled IPv6 for some reason, those apps would mysteriously no longer function. So the solution is reenable IPv6.

So anyway, many problems, as we know, for this rather painful round of updates. And of course it's no wonder, seeing this, that enterprise IT would be reluctant to jump onto updates immediately. You can imagine if enterprise-wide they had standardized on HP systems that all were blue screening after they restarted, and that caught them by surprise. So I'm liking the idea more and more of holding off for maybe a week at least on installing new features. And given the fact that 1803, which is the existing latest build, not the feature update, also suffered as a result of many of these problems in October, maybe just holding off on updates altogether. Or if nothing else, make an image of your system so that you're able to recover from something that might happen so that you're not in trouble until Microsoft gets around to fixing it.

Also, believe it or not, there's one more. Remember that on May 8th of this year Trend Micro, through their Zero Day Initiative, we discussed it at the time, notified Microsoft of a potentially serious remote code bug in their age-old and present everywhere JET database engine. It wasn't clear whether Microsoft didn't think this was that big a problem because it's not remotely exploitable. On the other hand, it is pervasive on Windows. Every version of Windows for decades has had the JET database engine in it. It was one of Microsoft's earliest database implementations, and it's still around for pervasive backward compatibility.

And remember that, as we discussed at the time, if you downloaded a piece of email that had a JET database engine file, essentially this is an interpreter buffer overflow-style exploit so that, if a user were to open that file, that could execute code of the attacker's choice on their system. And recall that prior to the September Patch Tuesday, Microsoft got back to Trend Micro a few days before and said, oh, we were unable to duplicate that, even though they had originally confirmed four months, 120 days earlier, that they acknowledged the receipt and said that they had verified that they had successfully reproduced the issue.

Well, Trend said sorry about that, but we gave you four months, so we're going to publish. So they missed the September patch update; right? So then on the 20th of September Trend went public with the news. And at the time we discussed it on the podcast, of course, and said that, okay, well, Microsoft barely missed it for September. Certainly they were going to get it for October.

Now we're at September 20th. There is that micropatch available from the Acros Security guys, which I think was a 22-byte patch that was an in-RAM patch, which the Acros technology, they have those things they call "micropatches" to fix little things like this quickly. And I said, you know, it's a third-party patch. Even though the patch itself is 22 bytes, that's a patch descriptor. You have to have a patching engine downloaded first. You have to sign up for an account with them. And it's like, let's not bother.

Okay, well, the other thing that happened for October is that Microsoft did implement a fix for this problem in the JET database engine, and they did it wrong. So the problem is still there. So now we have sort of a different concern, and that is probably for a month we're not going to have this thing properly fixed. And we know from a binary diff of the DLL that was fixed at the beginning of the month exactly what it is that Microsoft did. So this is a classic case of reverse engineerability. The guys at Acros Security who implemented the patch immediately for the original problem, their patch was broken by Microsoft's attempted fix of the JET database engine. They have a new patch - which is smaller, now it's only 18 bytes - which repatches Microsoft's breakage of their earlier patch.

But the issue is, for people who don't do this - and I still can't recommend a third-party patch, especially for something for which is there is no known public exploit happening. On the other hand, if it is publicly exploited, we may not know because it would be a targeted attack. It would be somebody knowing - and understand where we are at this point. At this point every single version of Windows in use has this remote code execution vulnerability in it today, fully patched. That is to say, every fully patched version of Windows has this problem such that a malformed file, which happens to be a JET database database, if executed, will run code on the victim's machine. So that file could be delivered through a web page download, through email, through whatever means; and, if executed, it runs code on that target machine.

And every fully patched Windows and every any patched Windows system now in use is vulnerable as a consequence of the fact that, first of all, Microsoft didn't fix it in four months. When they attempted to fix it in the fifth month, they did it wrong and left the vulnerability still there and exploitable. And we now have a month window, whereas before Trend didn't go public until the 20th of September, so it was only going to be a few weeks. Well, now we have four weeks. And if any bad guys started working on an exploit back when Trend went public, then they're rubbing their hands together because they just got a four-week extension on the availability of this thing being exploitable.

So we'll keep our eyes out for any sign of exploit. Microsoft says there's no known at the time. Microsoft, by the way, fully acknowledges the problem now, thinking that they had already fixed it. I've got a link in the show notes to their security guidance for the advisory, and this is CVE-2018-8423, where Microsoft says: "A remote code execution vulnerability exists in the Microsoft JET Database Engine." Yes, and still exists.

"An attacker," they write, "who successfully exploited this vulnerability could take control of an affected system. An attacker could then install programs; view, change, or delete data; or create new accounts with full user rights. Users whose accounts are configured to have fewer user rights on the system could be less impacted than users who operate with administrative user rights. To exploit the vulnerability," writes Microsoft, "a user must open/import a specially crafted Microsoft JET database engine file. In an email attack scenario, an attacker could exploit the vulnerability by sending a specially crafted file to the user, and then convince the user to open the file."

So Microsoft is writing that, believing that they had fixed it, which they didn't. And it is now every version of Windows in use has this presently available for exploit. So I will be frankly surprised if we're not discussing in a week or two that people are being attacked with this because it didn't get fixed.

So anyway, if anyone is concerned, 0patch.com. That is zero, the digit zero, P-A-T-C-H dotcom, is where Acros Security has their so-called "micropatches," which is in this case an out-of-bounds write, and the little patch script is 18 bytes long. And essentially, every time you boot the system, this reaches in and tweaks a few bytes in Windows to actually fix the problem, rather than Microsoft, that hasn't managed to fix it in five months. So anyway, maybe Microsoft will introduce an out-of-band patch, if it's found to be exploited. Again, targeted attacks. It does require user action and somehow to arrange to execute something. But it's certainly cause for some concern.

So Abdulrahman Al-Qabandi - I practiced that name.

Leo: Nicely done.

Steve: He's a Kuwaiti security researcher who discovered and published, after he was sure it was patched, a worrisome Microsoft Edge browser remote code execution bug. And of course this argues against the idea of waiting too long to apply security patches because this is worrisome because it is so trivial to exploit. So Microsoft did patch it this month. And it is a potentially potent and trivial-to-exploit, website-deliverable, remote code execution bug. So all that's necessary is some HTML and a bit of JavaScript, and to somehow induce the user of Edge on Windows 10 to press the Enter key.

So it's not clear, but it seems as though an advertisement could deliver this because it's just HTML and JavaScript, and we know that that's what ads are. And if something could get you to press Enter, then this runs on your system. Now, again, patched in October. So if there's a window of opportunity, it exists between the acknowledgement of this problem. He did wait to verify that the patch existed, that is, that Microsoft had fixed the thing that he had previously responsibly disclosed to Microsoft before posting, not only the details of the problem, but full working proof-of-concept code.

So the point is that we know that not all of the gremlins who inhabit the Internet are high-end exploit developers. One doesn't need to be one in order to leverage this particular exploit. I mean, it's laid out. So it's a little more worrisome than it would otherwise be that almost anybody could leverage this into a working exploit. So again, patched already. But for any systems that are delaying for whatever reason the application of October's patches, this thing is there, and it does allow remote code execution under the default browser, which is Edge, for Windows 10. So just sort of a heads-up, and another example of this very dynamic world that we're in right now with security updates and patches.

Yesterday, all four major browsers, and of course by implication their minor subsidiaries, announced in a coordinated announcement the planned deprecation of TLS v1.0 and 1.1. The WebKit blog was titled "Deprecation of Legacy TLS 1.0 and 1.1." Google called theirs "Modernizing Transport Security." Microsoft said, "Modernizing TLS Connections in Microsoft Edge and IE11." And Mozilla said, "Removing Old Versions of TLS." In the show notes I have a graph from the Mozilla announcement which I will explain in words, but it's just interesting to see it graphically, showing for the period of August through September of this year, so the most recent couple months, under the beta of Firefox 62, the relative ratio of connections by TLS version. So, okay. And we'll get to that in a second.

Backing up a little bit, TLS 1.0, which was, as we covered at the time, mostly just a renaming of SSL v3.0 - which is to say there really was no difference. They were the same protocol. It just became time to say, okay, we're not going to do Secure Sockets Layer anymore. We're now going to call this Transport Layer Security. It will enjoy its 20th birthday this coming January of 2019.

So at the beginning of next year, TLS v1.0 turns 20. And as such, and by any measure, it has been a spectacularly successful Internet protocol which, even today, nearly 20 years after its introduction, is without any flaws that are known to be serious enough to have forced its early retirement. But nevertheless, its age is beginning to show through its lack of support for 20 years of subsequent progress, mostly in its lack of support for modern security protocols, that is, the handshaking and encryption and authentication protocols which are negotiated by the client and server. There's been a lot of progress, not surprisingly, in 20 years that TLS v1.0 doesn't know about. So it's not that there's anything broken about it. It's just that the things that it's able to bring along in terms of its cipher suites, they're just not modern anymore.

So consequently, as I said yesterday, all four major browser vendors announced their coordinated deprecation, the end-of-lifing of TLS 1.0 and 1.1. Where we are today is that TLS v1.0 is down near 1%; 1.11% of all connections today are using TLS v1.0. And 1.1 never really got off the ground. It tweaked the protocols a little bit to solve some problems for which there were already other workarounds. So no one really bothered with it, such that it's currently at 0.09%, I mean, almost none. And anything that supports 1.1 also supports 1.0, so it didn't ever really need to happen.

The major protocol today is 1.2. And 93.12%, so better than 93% of all connections that were observed by Firefox during their period of monitoring were using 1.2. And already we've discussed a few months back the final ratification of TLS v1.3. That's at 5.68%. So 1.3 is where we're eventually headed. 1.2 is where we are solidly now, with only 1.11%. And let's see. In fact, if we add 1.0 plus 1.1, that brings us to 1.2% of those older protocols.

So when do they die? Even though they are barely seeing any use, they will not be formally killed until March of 2020. So, what, 15, 16 months from now, 17 months from now? So still plenty of time. But it is the case that the deprecation will be staged and made clear, I mean, so that everybody who is involved in this should know if they are in charge of websites which are not yet able to offer 1.2.

But it does mean that, when this finally happens in March of 2020, all the major browsers and their subsidiaries that are offshoots of their codebase will absolutely stop acknowledging TLS 1.0 and 1.1. And again, not really for any clear reason except that it's just time to move on. And it will allow the existing codebases to get cleaned up by no longer needing to back support a protocol that was originally created 20 years ago. So Google's Chrome will begin deprecating 1.0 and 1.1 in Chrome 72, such that visitors to sites that are unable to accept 1.2 or 1.3 connections, that is, only understanding 1.0 or 1.1, will begin to see deprecation warnings in the DevTool Consoles from Chrome 72 on.

So that's 72. Then there's 73, 74, 75, 76, all the way up to 81. It won't be until Chrome 81 that 1.0 and 1.1 will be completely disabled. So for what it's worth, you don't want to be anyone in charge of a website that cannot do TLS v1.2, or the Internet's going to seem very lonely after March of 2020. You won't be having many visitors to your site, certainly not using any contemporary web browser.

In happy news, Google has added a technology known as "Control-Flow Integrity" to beef up the security of their Android kernel, starting with the kernel which is present in Android 9, which is the most recently released Pie version. CFI, as it's known, Control-Flow Integrity, is an outgrowth of the LLVM project, which has been by all measures very, very successful. That's where Clang, Objective C, C, and C++ have implementations. It's a state-of-the-art compiler development project which began as a project in the University of Illinois, LLVM.

It probably originally stood for Low-Level Virtual Machine, I mean, that's what you would think the acronym would stand for. The project makes it very clear that it is not an

acronym; that their stuff has nothing to do with low-level virtual machines, although they did develop and have an intermediate representation. That is, one of the things that compilers do is they'll often compile from a high-level language, especially if they have multiple front ends, for example, if it can do C and Objective C and C++ and Clang, also known as C Lang, and potentially other languages, those front ends, the high-level end will all compile to a common intermediate representation which is then - and that's sort of like a midpoint stage or staging representation.

And then you have, depending upon which hardware architecture the compiler is targeted for, you then take that common intermediate representation and emit code for ARM or for x86 or X64 or whatever architecture you're targeting. So my guess is that that's where LLVM came from. They're distancing themselves from that as an acronym, for what it's worth.

Okay. So CFI is an outgrowth of that work. It's been around for a few years and has been already incorporated into other projects. Microsoft first incorporated it into Windows, and they called it "Control Flow Guard," CFG, which appeared in the Update 3 of Windows 8.1 back in November of 2014. And developers are able to add this, if they've got Visual Studio 2015 and later, by adding a /guard:cf flag to linker. So it's readily available and present. And as of Windows 10 Creators Update, which was the famous v1703, the Windows kernel itself is compiled with this Control Flow Guard.

So the central idea of both CFI and Windows implementation CFG is it's another approach to dealing with this Return-Oriented Program (ROP) exploitation. We've talked about this, when you lock the system down such that you can no longer execute your own code which you've provided in a buffer overflow when the buffer is on the stack because the stack has been marked as non-executable. Attackers, infinitely and always and ever clever, they said, oh, okay, well, we can't execute code we brought, so we're going to execute code that's already there.

And of course any OS kernel has got tons of code in it, waiting to be executed. And although none of the subroutines or functions that are there may do exactly what the attacker wants, what clever attackers have figured is that the ends of subroutines, like the last chunk of a subroutine, might do something useful, after which, as subroutines do, they return to the caller. So that's why this was known as Return-Oriented Programming, the idea being that you could knit together a malicious effect by just successively jumping into the tails of the subroutines in existing code, in order not to execute your own, but to execute code that's already there. But notice that what's happening is you are making a jump into somewhere that you weren't intended to go. That is, a subroutine is meant to be entered at the top and do its work and then exit at the bottom, not to be jumped into the middle of.

So anyway, the idea is to catch and prevent this malicious code reuse by checking the jump destinations of any data-driven indirect jumps. An indirect jump is one where existing code is jumping to another location where that location is specified by the data contained in a register. So it's indirect because the jump instruction itself doesn't say where to go to. The jump instruction looks at the contents of a register and uses that as the destination for its jump. That's called an indirect jump. So since bad guys might be able to change a register's contents, they are then able to execute a jump to the location of their choosing.

So Windows prevents this, believe it or not, by maintaining a bitmap of allowable valid jump destinations. And when a CFG application is being built, a bit of extra bitmap-checking code is added before every indirect jump instruction to verify that the location it is jumping to is a valid entry point, that is, something the system expects code to be jumping to, not code down in the middle of a subroutine which no valid code would ever jump to indirectly.

Chromium, the Chromium web browser, uses similar technology, the CFI technology under Linux, where the addition of CFI has been benchmarked to have an overhead, an execution overhead of less than 1%, but in this case it had a size overhead of about 15%. The benchmarker said, well, this hadn't been optimized for size, so don't hold that against us. But in today's plenty of memory computation world, the protection provided to a large attack surface such as a web browser probably merits the inflation of size by 15% if it really provides substantial protection against exploitation of the browser's own code because the browser, as we know, is the preeminent attack surface today.

So our news is that Google's Pixel 3 is the first Android device to ship with CFI protecting its OS kernel. To the degree that OEMs don't turn that off and sort of leave it alone and just customize around the OS, other non-Google devices will probably inherit this moving forward. And I expect it to become increasingly prevalent.

I did some poking around to see whether I could find any evidence of whether iOS was using CFI. They're of course heavy into Objective C, and they are using the LLVM project's work for iOS stuff. So I would be surprised if it wasn't in use. I found a research paper from several years back where they had looked at using a jailbroken iOS and shoehorning this CFI technology into iOS. But I didn't find any obvious indication that it is in use today. I'd be surprised if it wasn't.

You know, again, Windows started incorporating it four years ago. So it does look like it's yet another way of strengthening our software with, I want to say, "minimum overhead," certainly minimum execution time overhead, based on the Chromium browser benchmark. And certainly useful, if it really does, if it is effective in preventing this class of attacks, as it might well be.

Okay. So, oh, Leo. There was a report that was commissioned by Congress to ask the U.S. Government Accountability Office to take a look at the state of the Department of Defense's readiness against cyberattack. Now, we have to acknowledge that even the term "government accountability" is a bit of an oxymoron itself. But I have a link to a 50-page PDF. And I pulled some excerpts from it, one that just caught me by surprise that I thought was kind of fun.

So in their report the GAO says why GAO did this study: "The DOD plans to spend about $1.66 trillion to develop its current portfolio of major weapon systems. Potential adversaries," they write, "have developed advanced cyber-espionage and cyber-attack capabilities that target DOD systems."

What GAO found is sad. It won't surprise anyone. But, I mean, it's actually a little terrifying. "The Department of Defense faces," GAO writes, "mounting challenges in protecting its weapon systems from increasingly sophisticated cyber threats. This state is due to the computerized nature of weapon systems, DOD's late start in prioritizing weapon systems cybersecurity, and DOD's nascent understanding of how to develop more secure weapon systems."

Leo: Really? Nascent? It's just being born. Just now.

Steve: Apparently. They're just thinking, well, you know, maybe we should start. Maybe we ought to think about that, yeah. "DOD weapon systems are more software dependent," they write, "and more networked than ever before. Automation and connectivity are fundamental enablers of DOD's modern military capabilities. However, they make weapon systems more vulnerable to cyber attacks. Although GAO and others have warned of cyber risks for decades, until recently DOD did not prioritize weapon systems cybersecurity."

Let me read that again. "Until recently, DOD did not prioritize weapon systems cybersecurity. Finally," they said, "DOD is still determining how to best address weapon systems cybersecurity." So not only is it nascent, Leo. We have a committee, and the committee is going to figure out how we should consider going about addressing weapon systems cybersecurity. But wait, there's some examples here in a minute.

"In operational testing," GAO writes, "DOD routinely found mission-critical cyber vulnerabilities in systems that were under development, yet program officials GAO met with believed their systems were secure and discounted some test results as unrealistic. Using relatively simple tools and techniques, testers were able to take control of systems and largely operate undetected, due in part to basic issues such as poor password management and unencrypted communications." Okay. So this is like 20 years ago, right, in the commercial sector. And this is today.

"In addition, vulnerabilities that DOD is aware of likely represent a fraction of total vulnerabilities due to testing limitations. For example, not all programs have been tested, and tests do not reflect the full range of threats." And I wrote in my show notes here: "So begins a 50-page report detailing nearly total disregard for cybersecurity within the U.S. Department of Defense." And then something caught me short, and I grabbed this about halfway down under "Test Teams Easily Took Control," where they did detail a little more specifically.

They wrote: "Test teams were able to defeat weapon systems cybersecurity controls meant to keep adversaries from gaining unauthorized access to the systems. In one case, it took a two-person team just one hour to gain initial access to a weapon system, and one day to gain full control of the system they were testing. Some programs fared better than others. For example, one assessment found that the weapon system satisfactorily prevented unauthorized access by remote users, but not insiders or near-siders. Once they gained initial access, test teams were often able to move throughout a system, escalating their privileges until they had taken full or partial control of a system.

"In one case, the test team took control of the operators' terminals. They could see, in real-time, what the operators were seeing on their screens and could manipulate the system. They were able to disrupt the system and observe how the operators responded. Another test team reported that they caused" - get this, Leo - "a pop-up message to appear on users' terminals instructing them to insert two quarters to continue operating."

Leo: Game over, man.

Steve: "Test reports indicated that test teams used simple tools and techniques to disrupt or access and take control of weapon systems."

Leo: This is terrifying.

Steve: "For example, in some cases, simply scanning a system caused parts of the system to shut down. One test had to be stopped due to safety concerns after the test team scanned the system." Yes. In other words, as horribly insecure as our plastic blue-box SOHO routers are, they are arguably more robust than major United States Department of Defense weapons systems. Our tax dollars hard at work.

Then, finally, I'll just finish with another quote from this report: "Program offices were aware of some of the weapon system vulnerabilities that test teams exploited because they had been identified in previous cybersecurity assessments. For example, one test

report indicated that only one of 20 known cyber vulnerabilities identified by a previous assessment had been corrected. The test team exploited the same vulnerabilities to gain control of the system. When asked why vulnerabilities had not been addressed, program officials said they had identified a solution, but for some reason it had not been implemented."

Leo: Oh, my god.

Steve: What do you know. "They attributed it to contractor error." That's right, point your finger.

Leo: It's a big iceberg. We only - it's a big iceberg.

Steve: Oh. "Another test report indicated that the test team exploited 10 vulnerabilities that had been identified in previous assessments." In other words, yes, we have problems. And we knew it, and we've identified a solution, but for some reason it hasn't been implemented. Wow.

Leo: Well, at least this report exists. I mean, I hope they listen to it and do something about it.

Steve: Hope they turn the temperature up, yup. Our friends at the EFF are warning of the widespread use or, well, increasing use of facial recognition. And I'm glad they're on the lookout for us. They said: "A proposed amendment to the Chicago municipal code would allow businesses to use face surveillance systems that could invade biometric and location privacy and violate a pioneering statewide privacy law adopted by Illinois a decade ago."

So again, this would be a City of Chicago municipal code which essentially is attempting to override a statewide good privacy law that was adopted a decade ago. The EFF joined in a letter with several allied privacy organizations explaining their concerns, which include issues with both the proposed law and the invasive technology it would irresponsibly expand, so they say.

And so to give a little bit of background they said: "At its core, facial recognition technology is an extraordinary menace" - this is the EFF writing - "to our digital liberties. Unchecked, the expanded proliferation of surveillance cameras, coupled with constant improvements in facial recognition technology, can create a surveillance infrastructure that the government and big companies can use to track everywhere we go in public places, including who we're with and what we're doing.

"This system," they write, "will deter law-abiding people from exercising their First Amendment rights in public places. Given continued inaccuracies in facial recognition systems, many people will be falsely identified as dangerous or wanted on warrants, which will subject them to unwanted - and often dangerous - interactions with law enforcement." They say: "This system will disparately burden people of color, who suffer a higher false positive rate due to additional flaws in these emerging systems. In short," they write, "police should not be using facial recognition technology at all, nor should businesses that wire their surveillance cameras into police spying networks."

And then they said: "Moreover, the Chicago ordinance would violate the Illinois Biometric Information Privacy Act (BIPA). This state law adopted by Illinois statewide in 2008 is," they write, "a groundbreaking measure that set a national standard. It requires companies to gain informed, opt-in consent from any individual before collecting biometric information from that person, or disclosing it to a third party. It also requires companies to store biometric information securely and sets a three-year limit on retaining information before it must be deleted. It empowers individuals whose rights are violated to enforce its provisions in court."

So they conclude, well, actually it goes much longer, but I concluded quoting them, saying: "Having overcome several previous attempts to rescind or water down its requirements at the state level, BIPA now faces a new threat from this recently proposed municipal amendment in Chicago. The proposal to add a section on 'Face Geometry Data' to the city's municipal code would allow businesses to use controversial and discriminatory face surveillance systems pursuant to licensing agreements with the Chicago Police Department." Anyway, it goes on, but we get the idea.

And so once again we have a situation where the technological advancement and progress which we all herald is creating new capabilities that our established legal frameworks and assumptions have not been updated yet to address. And of course we all know because it was just recent when Apple's iPhone X introduced their Face ID, a great deal of concern was raised about its privacy implications, and Apple went out of its way to explain how it was purely a local recognition and unlock capability and that this wasn't giving Apple access to our 3G face geometry. It was entirely contained within the phone, very much like the finger ID beforehand. So anyway, I'm very glad that the EFF is watching our backs. And after I wrote that, I thought, yes, and not our faces. So, yay.

In my browsing around for things to talk about that I thought were interesting, I ran across in Bleeping Computer a report from DuckDuckGo. They publish the activity reports just in numbers of how much use they're seeing. And I put a chart in the show notes showing essentially exponential growth, I mean, a rapidly accelerating growth in the use of DuckDuckGo. You can see the numbers at DuckDuckGo.com/traffic because they publish this traffic report. And they proudly claim, I mean, they proudly pronounce themselves as the search engine that doesn't track you.

What annoys me, and I've mentioned this before, is that Google's search links are now all redirecting, which is annoying because I want to be able to right-click and do a "save as" in order to grab the link that Google's search is referring me to, and they're all crazy nonce-burdened Google redirect links. Which means, as we know, that Google, who knows who I am because there's my face on the browser page, I'm logged in with Google stuff, they know everywhere I go. Every link I click on Google search is being tracked.

So I was curious, and I brought up a search under DuckDuckGo.com. And sure enough, every link there is a direct link to the site that you're searching for. So I know that there's a chunk of our listeners who are interested about their privacy and are annoyed by tracking. So I just wanted to put DuckDuckGo on everyone's radar, if it wasn't already. Oh, and I got a kick out of this, Leo. Under donttrack.us, D-O-N-T-T-R-A-C-K dot U-S, the second slide there is kind of fun. It shows Google preventing ads which are tracking. Clearly, DuckDuckGo is presenting themselves as an alternative to Google.

Anyway, it's refreshing to see an alternative. And in fact I think it was in the Bleeping Computer report, there were a large number of people that were responding to it, saying, yeah, it's all I use. It's the engine I've chosen and so forth and so on. So it looks like it is in fact very popular and, as we can see from the chart, growing very rapidly.

Ten weeks from now, at the end of 2018 - wow, 10 weeks from now, whew, 2019 - PHP's 5.x branch support ends. I'm very glad that I've chosen PHP 7. I mean, and I kind of had

to work at it. When I set up GRC's SQRL forums, I'm using XenForo, which is a PHP-based web forum solution. And it works fine under 7, but all the defaults were for 5. And that's still the case, broadly. I mean, even though 7 has been around for quite a while now, for example, WordPress has their minimum requirement at 5.2. Joomla is at 5.3.

And as I said, I had to kind of dig a little deeper and push and work in order to use 7.2 as my own platform moving forward. I'm glad I did because, I mean, I could change, but it's easier not to. Which is exactly the point here. Right now 78.9%, 78.9, so just shy of four out of every five Internet sites, are PHP powered. That is, to varying degrees they run on PHP. On the stroke of midnight, when support for the 5.6 branch of PHP ends, at that point 62% of all Internet sites, which is that percentage of the total still running PHP 5.x, will stop receiving security updates for their server and their website's underlying technology, which at that point will, moving forward, expose hundreds of millions of websites, if not more, to potentially serious security risks.

Now, backing away from that, PHP is pretty secure. It is relatively stable. And it's been a while since we've seen a major problem with it. But when you consider, I mean, essentially this increases the target-rich environment that attackers use. It is open source, so the code can be scrutinized. Starting with the beginning of 2018, 62% of all Internet sites that don't move will be vulnerable if a problem is found, and that version will not be patched.

So anyway, I just sort of wanted to put this on the "get ready for next year" radar because I wouldn't be surprised if, when it's known that vulnerable sites can't easily update, it may paint a bigger target on those sites that don't. And for what it's worth, as a heads-up to everyone, 10 weeks from now, no more 5.6 branch updates. So it might be worth considering, if you have an actively functioning operating PHP-based site, as apparently four out of every five sites on the Internet are, that you consider biting the bullet and making the move over to the 7.2 branch, which is where we are with the v7 series of PHP. Which is not to say that there'll be a problem. It's got a good security record now. It's settled down and is being well run. Let's hope that there isn't a big problem that appears.

And I mentioned skepticism about AV protection. Which AV is the best, or the answer to that question, is a consequence of who you ask. There is a site that to me looked unbiased and comprehensive, known as AV-comparatives.org. Again, I'm not making any representations for it. They explain, they said: "Our Real-World Protection Test is currently the most comprehensive and complex test available, using a large number of test cases. Currently," they write, "we're running this test under Microsoft Windows 10 Pro RS4 64-bit with up-to-date third-party software such as Adobe Flash, Adobe Acrobat Reader, Java, et cetera. Due to this, finding in-the-field working exploits and running malware is much more challenging than, for example, under a non-up-to-date system with unpatched vulnerable third-party applications." So they're testing against the sort of systems that a responsible person who, well, hasn't removed Adobe Flash is using.

So in the chart that I have in the show notes, they said: "The results are based on the test set of 197 live test cases where malicious URLs were found in the field, consisting of working exploits, for example, drive-by downloads and URLs pointing directly to malware. Thus exactly the same infection vectors are used as a typical user would encounter and experience in everyday life. The test cases used cover a wide range of current malicious sites and provide insights into the protection given by the various products, using all their protection features while surfing the web."

So a couple things. First of all, the chart shows essentially the percentage of among those 197 live test cases, by AV, how many of them were found. Avast, AVG, BitDefender, F-Secure, Microsoft, Trend Micro, and is it Vipre? - yeah - were the only ones with perfect scores. Of those, Microsoft had by far the greatest number of false

positives, with Trend Micro in second place, F-Secure in third place, and the others lower. And you do see a correlation here. There is a tendency for those which were providing the best protection to have the highest false positive rate. So that suggests that they're more heuristic. They're cutting, they're slicing this a little further in favor of detection rather than not.

What I like is, if the takeaway is - as you and I, Leo, have been saying now for a while. You know, given that add-on AV is having a problem under Chrome such that it's being essentially kicked out eventually, and that it represents an attack surface which can itself lower an individual's security, and given that Microsoft is one of the six AVs with a perfect score, maybe just using the system built into Windows is sufficient protection, downside being some possibility of false positives. But I'd rather that than have things missed. So anyway, I just thought this chart was interesting, and it does support the contention that running with what Windows now provides us is providing protection that is as robust, to the degree we believe this test, as the other contenders.

Okay. So a little bit of "closing the loop" feedback. I saw a tweet from - and I didn't know how to pronounce his name. I thanked him for the tweet, and he wrote back saying his name is - he's actually Johan. He said: "Hi. Listener of your show for about a year now, maybe longer. What are your thoughts on OCSP? Had an issue the other day where I could not access my website via Firefox or Edge; could in Chrome. Turned off OCSP validation in Firefox, and it worked." He says: "Keep the setting off or turn back on? Thanks."

So as we know, OCSP is the real-time certificate validation which is a way of solving the revocation problem. That is, on the fly, a web browser can be asked to check for revocation. One of the problems is that most web browsers have defaulted to defaulting open. That is, if an OCSP server does not respond, then they'll just say, oh, well, at least we don't know that the certificate we're being asked to trust has been revoked. We don't have an answer one way or the other. Because the certificate itself provides the URL for the browser to use to check on its real-time revocation status.

So to answer Johan's question, I would, I mean, certainly running with it enabled as he had it in Firefox is more secure. There is the problem that OCSP servers are sometimes not up. So if you fail closed, that is, failed to not trust a certificate that you haven't been able to affirmatively verify is still valid, then you're going to have this problem for some sites.

I would say, if it's not too much trouble, turn it off to go somewhere that you trust. Be very careful, if you turned it off to go somewhere, that you really are there because, I mean, this would be exactly the attack scenario, right, is that there would be no CSP support for that cert, and it would have been revoked. And so somehow a bad guy would be blocking your access to it in order that you didn't know that it was broken. So be very careful while you're there. And once you're done, turn it back on again. OCSP is much more robust today than it was 10 years ago. It's only going to get better moving forward. And I think ultimately we're going to get to a point where we can trust it.

Stapling is what we really want. Stapling is the right solution. I went over to GRC just recently and noted, I mean, I went over to SSL Labs and looked at GRC relative to cipher suite support. And I smiled when I saw that the server I'm using supports OCSP stapling, and it's in use. What that means is that the web server itself goes and gets a fresh OCSP attestation and includes it with the certificate that it provides the web browser. So the web browser doesn't need to go out and separately get the OCSP certificate. So it's really, I mean, as I've discussed where we've talked about revocation, OCSP stapling is absolutely the right solution. And as servers evolve to support it, that solution will be increasingly available. And then OCSP won't be a problem.

Also Timo Gruen tweeted: "On Bloomberg and Supermicro," he says, "why put the chips ON the boards where they can be found? Modern motherboards have so many layers, and the chips are so small, you could easily sandwich them INTO the board. Good luck finding that." And as a matter of fact, that was noted in the article. There was some comment among experts who noted that, with motherboards being multilayer, you absolutely could actually bind the chip into the layers of the motherboard because a lot, I mean, the chip itself is vanishingly small. It's the packaging which supports the pinouts on the chip which gives it any substantial size. And even then, as we know, that's not much. So I just wanted to acknowledge that that was something that was mentioned in the Bloomberg article that I failed to mention in last week's podcast.

Someone tweeting as the "Sultan of Saki," whose actual name he signed off "Josh Fenton," he said: "Hi Steve. A friend recently asked me how they should go about selling some used hard drives on eBay. I explained to him the importance of wiping them, and as I was doing so I started thinking about how drives can swap out bad sectors for spares. It occurred to me that it is possible that, if the drive was ever used unencrypted, that any data in sectors that had been swapped out would therefore be completely inaccessible to the OS or application layers, thus making it possible to wipe the data they contain. If my analysis of this issue is correct, then this is yet another reason for users to always ensure that they encrypt their drives immediately upon installation; and in no case should they copy data to the drive until the encryption is complete. Do you think my thinking is reasonable here, or am I missing something? Thanks, Josh."

And, no, you are correct, Josh. And this of course applies equally to thumb drives and SSDs also, which have the ability to maintain a reserve pool of storage space and swap it out. There is in the latest spec, which I will be programming to when I return to SpinRite 6.x work, there is a secure wipe feature. And I'll be looking at it closely to see whether the manufacturers have uniformly supported that as we would hope they have and allow the secure wipe to also wipe sectors which have been spared out, in which case that might be like I've talked about how I intend to produce a product called Beyond Recall, which will itself take responsibility for doing an extremely fast, but also secure wipe. And it may incorporate the secure wipe feature as an option, depending upon what we learn about that when we look more closely.

But it is absolutely the case that one way to avoid the problem is by using whole drive encryption, either as an add-on like BitLocker, or down at the BIOS. As we know, all drives now have the ability to support a password. When you give the drive a password that the BIOS uses, then that drive is also doing whole drive encryption of itself. You would then want to make sure you remove the password. But in the process of removing the drive from the motherboard, then it will be passwordless, and the drive will be secure. So that is another way to operate is to take the trouble to put a password at the BIOS level.

There are problems with doing that. You need to be careful about that. You need to, for example, if you're going to remove the drive for use outside, you need to remove the password from the drive in order to make the drive data accessible without the motherboard BIOS to provide the password. So you have to use that with some caution, but it can be done.

Gary Napier asked: "Hey, Steve. Thanks for all the great info. Do you know of any way to check your router to see if you are infected with VPNFilter?" And I would say, first of all, we know that it exists persistently. That we know. It was VPNFilter which was found and now has, what, six different means of providing persistence for itself. So rebooting the router won't help you. What I would do, if you have any reason to suspect that the router model is one of those which is known to be vulnerable, is I would reflash the BIOS, I mean, reflash its firmware using the latest one from its manufacturer, and then immediately lock it down.

So rather than, I mean, there's probably no good way without really digging in and looking into the file system. And even then, it's going to try to hide itself. I would just go back to the factory. Go back to factory settings, reflash the router, and then lock it down carefully. And of course reflashing with the latest firmware may solve any known vulnerabilities at that time. But certainly then, when I say "lock it down," I mean close off any publicly accessible services that you're not actively needing.

Okay. And I just did want to finish with a tweet from John McAfee, for anyone who might have been worried after seeing this and experiencing the "presidential alert" message. Our friend John McAfee tweeted "The presidential alerts: They are capable," he tweets, "of accessing the E911 chip in your phones, giving them full access to your" - you know, "them" - "full access to your location, microphone, camera, and every function of your phone. This is not a rant. This is from me, still one of the leading cybersecurity experts. Wake up, people," tweets John.

Leo: Wake up. Wake up, people.

Steve: Wake up, people. Anyway, for anyone who might be concerned, there is no E911 chip. There is a 911 function which does require that our phones provide location information for the 911 service to locate us. But that's the limit of it, and it's got nothing to do with presidential alerts. The presidential alert, as Leo, you said correctly when we were discussing this at the top of the podcast before we began recording, is a broadcast. And it doesn't suddenly peg the instantaneous location of every U.S. citizen, all 300-plus million of us.

Leo: That'd be quite the thing, yeah.

Steve: With cell phones. So, yeah.

Leo: We know where you all are. Don't move.

Steve: Thank you for the public service announcement, John, but not a problem.

Leo: Yeah, the government wants to know where we all are. I guess it could be used to identify just where I am. Right here.

Steve: So, Good Samaritans. So first of all, we know routers are a problem. Exactly how severe is the problem? A study was conducted recently by the American Consumer Institute, ACI, a U.S. consumer nonprofit. They found that five out of six home routers are inadequately updated for security flaws, leaving the devices and indirectly their users vulnerable to hacking. Not surprisingly, but still five out of six. The study, and I have a PDF link in the show notes, analyzed a sample - a relatively small sample, so that might make the study a little questionable - but 186 SOHO [Small Office Home Office] WiFi routers from 14 different vendors. So they weren't looking for population as much as here's 186 different makes and models of routers in use. What's their status?

So 14 different vendors. They looked at the firmware version the routers were running and searched public vulnerabilities databases for known security flaws affecting each device's firmware. So again, we don't know that that's remotely executable. We just

know that there's a security flaw that's known in the firmware which had not been patched for any one of these 186 routers. They wrote: "In total, there were a staggering 32,003 known vulnerabilities found in that sample."

So actually I guess the small sample size helps to keep this under control: 186 routers from 14 different vendors found to collectively have 32,003 known vulnerabilities. They wrote: "Our analysis shows that, of the 186 sampled routers, 155" - which is 83%, so a little bit more than four out of five - "were found to have vulnerabilities to potential cyberattacks in the router firmware, with an average of 172 vulnerabilities per router, or 186 vulnerabilities per router for the identified 155 routers," that is, of the 83%. "Of the total 32,003 security flaws, more than a quarter were vulnerabilities that received the two highest severity ratings of 'critical' and 'high-risk.'"

Okay. So three quarters of them weren't even critical or high-risk. So they were a technical vulnerability, but nothing you need to worry about critically. Still, a quarter of 32,000 is, what, 8,000. So still significant. They said: "Our analysis shows that, on average, routers contained 12 critical vulnerabilities" - of course you only need one - "12 critical vulnerabilities and 36 high-risk vulnerabilities across the entire sample." So a significant problem.

Last week news surfaced of a mysterious, as I said at the top of the show - I recall from my reading into this that he was Russian speaking - a mysterious vigilante grey hat hacker who is patching people's outdated MikroTik routers. From what I read, he is not performing a remote firmware update, which, okay, well, he's not. Whether we think that would be better or not is up in the air.

So as I also mentioned, I've noted that my advice has raised some controversy, that is, my advice being any device which a naive user can use that creates an exposure for them, and even for others, for example, in its use as a botnet host which is creating massive floods which are increasingly difficult to deal with, should have manufacturing-managed automatic update somehow. People have pushed back, saying wait a minute, I hate the idea of my router updating itself. So, okay. I get that. I mean, I understand that.

So how about when you use the router the first time, you are asked if you want automatic updates. Or you're told that, unless you turn it off, they are on. I think that's the right way. I think it needs to default to maintaining itself, to phoning home, checking for an update. Now, maybe it only notifies rather than performs the update autonomously. That would be, again, another option, a back step, if it's something that can notify. A light bulb can't notify unless it refuses to turn on or flickers or does something. But then the typical user is going to have no idea what's going on.

On the other hand, a light bulb probably can't actually have firmware updates. But hubs that run IoT devices, or certainly our routers can. And so they could, for example, redirect their user to a notification page, saying, "There is a known high-risk vulnerability in this version of the router's firmware. Sorry about that. Please update the firmware. Press this button to do so."

So again, there are many compromises that could be made. I obviously have no problem with having a sophisticated user turning those things off. And as we've discussed in the case of MikroTik over the last few months, MikroTik kind of has options for allowing the router to be profiled. But even the least sophisticated profile doesn't protect the user who says, "I don't know what I'm doing, so I'm choosing this profile." Even that isn't configured by default to protect them.

So these defaults have to change. And it is entirely because of the default settings that MikroTik is in the trouble that it is today. So now we have the case of a vigilante breaking

the law because it is absolutely, definitely, unequivocally against the law to hack somebody else's router without their knowledge or permission, even if it's with the best of intents and to help them. And note also that what this guy is doing is, as I understand it, not updating firmware, but bringing their firewall to bear, closing ports which are open, but it's possible that in some cases those are open deliberately. That is, unfortunately, they are open by default. They should absolutely not be open by default. In that case, if they were open, they would be open on purpose, in which case closing them would probably break something.

So the problem we have is that with it being open by default, a vigilante, as well intended as they may be, is unable to determine whether or not in fact they are in use by somebody who needs remote access to a router, such that bringing up firewall rules to close them will break functionality which is needed. So it's a mess. Anyway, I just - it got a lot of coverage in the press. And Leo, what do you think about the idea of somebody coming along and fixing things?

Leo: It's terrible. We've seen this before. The chance of him doing something inadvertently bad are high. And it is highly illegal. And, no, it's not okay in any respect.

Steve: Yup.

Leo: No. How could you justify this at all?

Steve: Yup. You're right. I mean, there are too many ways that it can go wrong.

Leo: Now, what if, however, MikroTik pushed out a firmware update? That's fine. So Mr. Russian Guy, go to MikroTik and help. I guess MikroTik doesn't have an automatic update facility, do they.

Steve: No.

Leo: Now, what if they took advantage of the flaw to do it? No, that would be wrong.

Steve: Yeah. MikroTik can't. Now, if they're...

Leo: Okay, I got one. What if they popped up a message somehow? What if they used this technique to at least alert people that there was a flaw and that they could go get firmware updates? How about that?

Steve: Wait. MikroTik or the Russki?

Leo: Not the Russki. MikroTik.

Steve: Okay, okay. I guess it's a function of what the fine print of their license says. I mean, there are MikroTik routers that have auto-update features. So they're not completely naive to the idea of updating their routers. They just don't have it on by default. However, I should mention that I'm beginning to see reports from people who are saying that the updated firmware in various consumer routers is starting to offer auto-update as an option.

Leo: Oh, I see that all the time, yeah.

Steve: Yes. So yay for that. And I don't have to tell our users that, yes, turn that on. You want that. And the problem is right now none of our routers - I'm not aware of a single router that notifies its user proactively. You log into the router. And, you know, for example, in the case of a Netgear, I have a couple Netgear routers, there's a flashing exclamation point saying check for firmware updates. And it's like, oh, well, it would have been nice to be notified.

Leo: Yeah, yeah.

Steve: But we're not being notified. So I agree with you, Leo. If there was, I mean, if we gave them permission to intercept, to do a browser page intercept and put up a page saying you've got obsolete firmware, and it's obsolete in a bad way. That would be a service. I have a hard time imagining somebody being upset by that. But I should mention there are reports of people being infuriated that some random Russian guy has changed their firewall rules, as you can well imagine.

Leo: Yeah, yeah. No, no, this is never okay. And don't think it is, kids. Knock it off. Well, Mr. G., I think we've come the end of this edition of...

Steve: So it's Bad Samaritan, not Good Samaritan.

Leo: Not Good Samaritan. I don't think there's any conceivable case it's a good idea.

Steve: No.

Leo: No.

Steve: No. At the same time, it's nice to have 100,000 fewer vulnerable MikroTik routers.

Leo: Maybe. Maybe they're not vulnerable. We don't know what the guy did.

Steve: True.

Leo: Maybe they're more vulnerable.

Steve: True.

Leo: Right?

Steve: True.

Leo: Steve, you're always a breath of hot - of fresh air. I'm just teasing you. It's always a pleasure. And I know people listen, wait all week to listen on Tuesday, 1:30 Pacific, 4:30 Eastern, 20:30 UTC. They wait all week for their Security Now! fix. And actually, you know, if you don't have time to do it live, and that livestream is at TWiT.tv/live, you can get downloaded editions. Steve has audio and transcripts you can download and read along as you listen at GRC.com. That's where SpinRite lives, his bread-and-butter, world's best hard drive maintenance and recovery utility, plus a lot of other free stuff Steve gives away. Lots of information. GRC.com. He's on Twitter, @SGgrc. That's a good place to keep up with him. But also, if you wish to communicate with him, you can do it there via direct message. He accepts direct messages from anybody.

Steve: I do, I do.

Leo: Although that's crazy. But he does. You can also go to GRC.com/feedback and leave a message there. You can get audio and video files from us, TWiT.tv/sn. That's where the Security Now! files are, TWiT.tv/sn. Or subscribe in your favorite podcast application. Details on how to do that, TWiT.tv/subscribe. You know, if you've got one, you know how to subscribe. Just search for Security Now!, and that way you'll get it the minute it's available. Again, audio or video. Although as Steve says, and I agree, who the heck wants to see us? But we show your Picture of the Week and stuff like that. There's stuff to see here.

Steve: Yes.

Leo: Steve, have a great week.

Steve: Will do, my friend.

Leo: See you next time.

Steve: Right-o. Bye.