Ripple20 Too

Description: This week we look at news in the shortening of certificate lifetime change, at Apple's decision to deliberately ignore support for a bunch of new Web APIs, at Apple's announcement of DoH support, at some troubling Mozilla/Comcast news, at some welcome legislation to head off the use of facial recognition, and at another less welcome attempt to outlaw strong encryption. We also look at the growing legislation against mandatory "chipping" and remind our listeners about the utility of VirusTotal. Then, after catching up with a bit of miscellany and listener feedback, we revisit last week's very worrisome revelation of the many flaws in a very widely used embedded TCP/IP stack. There's much news there.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-773.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-773-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Lots to talk about, including the Face Recognition Moratorium Act and, oh, boy, the new encryption law about to be voted on in Congress. He'll also talk a little bit about Ripple20 once again. It's getting even worse. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 773, recorded Tuesday, June 30th, 2020: Ripple20 Too.

It's time for Security Now!, the show where we cover your security, your privacy, and your computer and how it works, all with this guy right here. He's the king of Security Now!, the guy in charge, Mr. Steve Gibson. Hello, Steve.

Steve Gibson: Leo, great to be with you again. Yes, I'm saluting with my other hand because the microphone is blocking my normal salutation. Outside the house is a wood chipper, which is grinding up lord knows what, sounds like bricks.

Leo: Uh-oh.

Steve: But the good news is these amazing Heil microphones we use are extremely directional. And so even though it's kind of a pain due to my setup to put the microphone on the other side, I had sufficient notice to do so today, and the microphone is aimed away from the wood chipper, rather than right at it. So I figured that was a good thing to do for the sake of our listeners, whom we care so much about. That's why we're here every week for 773 weeks. That's our episode for the end of the month, end of June.

More news - details, actually - have some to light about the really horrific Ripple20 vulnerabilities. And so with a little tongue in cheek I called this episode "Ripple20 Too," T-O-O, because now that we know a lot more detail, you know, exactly which hundreds of millions of devices are vulnerable to this longstanding, multidecade, very worrisome TCP/IP stack vulnerability set, I thought we needed to follow up with that. And there wasn't anything else that was breathtaking that happened, although lots of interesting news.

We've got the actual shortening of certificate lifetime change, which we'll remember was something that Apple surprised everybody with back in February. Other people are coming onboard. So we need to update on that. We also have Apple's interesting decision to deliberately ignore their support for a bunch of new Web APIs. They've said no, thank you. And frankly, I can't disagree.

Leo: Yeah, when I read them, I thought, the browser could tell you that? I was like, wow.

Steve: I know. And that's, you know, that's what's so nutty about what the W3C Consortium are doing. They've got to be sitting around in some sort of brainstorming meeting, asking themselves, what else can we do?

Leo: What other features can we build in?

Steve: Yeah, yeah, like how many icons are on the user's desk. It's like, what? Well, that could be useful. Oh, okay, let's put that in there. Anyway, we've also got Apple's kind of weird announcement of the way they're going to be implementing DoH and, yeah, DOH and DOT, which we're going to cover. Some troubling, believe it or not, Mozilla/Comcast news.

Leo: Uh-oh.

Steve: Also, we've got some welcome legislation to head off the use of facial recognition; a less welcome attempt to outlaw strong encryption, another one, separate from the EARN IT Act, but with a common sponsor, not surprisingly.

Leo: And even more clear in its desire. Oh, boy.

Steve: Yes, it is very explicit, yeah.

Leo: Yeah, yeah.

Steve: It doesn't have the slime factor that EARN IT did, but it's still, like, in your face. We're also going to look at the growing legislation against, thank goodness, mandatory chipping of people, not pets, but yes, people. And I wanted to remind our listeners about the utility of VirusTotal as a consequence of my having touched it recently, and also some news there. Then we're going to catch up with some bit of miscellany, some listener feedback, and then revisit last week's very worrisome revelation about the many

flaws in a, as we know, very widely used embedded TCP/IP stack. And I still have the feeling, even with all this, that we're just looking at the tip of the iceberg. But we will see.

Leo: Bad news. I think Anthony Fauci says that you should stop touching your VirusTotal, but we'll talk about that later; okay?

Steve: Good idea. So last February we talked about Apple's surprise announcement during the CA Browser Forum that, in the future, and this was just unilateral, that in the future it, for Safari, on all of its platforms, would be rejecting any web server certificate having a not-valid-before date, which is technically the way the date range is stated, so it's very clear, not valid before date after August 31st of this year, of 2020, and which has a certificate lifetime greater than - total lifetime greater than 398 days. So in other words, starting just two months from now, that is to say, from September 1st on, all CA certificates issued for use by web browsers must be issued with a one-year plus 33 days lifetime or shorter, not longer.

So this is the death of the more arguably convenient two- or three-year web server certs that we've traditionally been using. Essentially, sort of Apple biting the bullet and pushing an issue that the various non-certificate authority participants in the so-called CAB Forum, the CA Browser Forum, had been asking for for a long time. Google had put forth this issue, this measure for a vote at the prior meeting, and it had been voted down in a partisan vote by the certificate authorities that said, no, we don't want to shorten server certificates to a year.

Well, Apple said, okay, tough. We're just not going to accept any. And arguably, Safari is strong enough that basically they forced the issue. So when I talked about this initially in February, I discussed the many implications of this in great depth and detail. So I'm not going to go into all that again. If anyone has joined us since then or wants a refresher, it's back in February.

The reason this is back in the news is that now the other two significant browsers in the industry, Mozilla and Google, plus all Chromium-based offshoots of Google's Chrome browser, have also announced their exactly aligned policies.

Leo: Yes. That's really interesting, wow. That'll hurry this along.

Steve: Google's Ryan Sleevi, yeah, he posted as sort of like their equivalent of things we're going to change, in the Chromium blog he said: "Enforce 398-day validity for certificates issued on or after 2020-09-01," September 1st of this year. And then the body of the message is: "Enforce publicly trusted TLS server certificates have a lifetime of 398 days or less, if they are issued on or after September 1st, 2020." And he said: "Certificates that violate this will be rejected with" and then the error is "ERR_CERT_VALIDITY_TOO_LONG and will be treated as mis-issued."

And also, following up, Mozilla's Kathleen Wilson posted: "Limit reuse of domain name verification to 395 days," and that was #206. And I think she did say 395 because I'm sure I copied and pasted. So they're off by three, it should be 398, I believe, because I remember one year plus 33. And that's to sort of give people a little bit of fudge room.

So there is a long and very interesting discussion, for people who like such things, among the industry insiders who are the ones who make these essentially earth-moving decisions. So I've included the Google Group's discussion thread link in the show notes

for anyone who's interested. I mean, it's back and forth and a lot of discussion. But basically it comes down to, well, you know, this is what we wanted. Thank you, Apple, for biting the bullet. We're all going to jump onboard.

And, you know, the certificate authorities will end up changing their model. Rather than, for example, you having to have a cash transaction annually, you'll be able to purchase some block of time that you want to have certificates from them. And I imagine, since that does create a little bit of lock-in, that they may extend that. They may say, hey, stay with us, commit to staying with us for 10 years, and we will lower the per-year cost of certificates. And then it'll be like, you log into your account and basically you reissue a certificate before the one you have expires.

What this will also do, you know, we always run across instances where people are forgetting or letting it lapse, or maybe it's a holiday, or it's a COVID-19 event, one way or the other server certificates are expiring, and they're finding out only when people are screaming that they can no longer access the website. So maybe it being an annual thing, as opposed to for example every three years, which maybe you're more likely to forget, that might help prevent that. And I know in my relationship with DigiCert I'm getting email from them all the time saying, hey, this or that's getting ready to expire, so don't let that slip by. So you certainly want to be, I mean, in general, absolutely want to make sure you have an email address that allows you to receive important security notifications.

I would argue that notes from your certificate authority is one class. And as we'll be seeing in several of the stories that we have this week, it's absolutely important, crucial for the software that you're using, the vendors of the software that you're using have a way to reliably get a hold of you to tell you when there's some action you need to take. And when you get that kind of notice, at least open the mail. Judge for yourself how important it is because it can be really crucial.

We were talking at the top of the show, Leo, because you were looking also at these Web APIs, I mentioned that Safari has chosen to eschew adopting these 16 Web APIs for the sake of user privacy. I would argue, who cares? Like the magnetometer? Okay. Apple has decided not to build these newly defined APIs into Safari because they both individually and collectively could pose a credible threat to user privacy by opening new opportunities for user fingerprinting. So let's quickly enumerate these.

Web Bluetooth, which allows websites to connect to nearby Bluetooth LE devices. That one, I don't know. I could see some uses for that, frankly, as the designer of SQRL. If we reliably had Bluetooth connectivity to the browser, that would create a direct connection between somebody's phone and the browser to create a local authentication loop to close that loop that no bad guy in Russia or China or some other foreign land could intercept. So that one's a little near and dear to my heart. But it's not available in Chrome.

The MIDI API, right, M-I-D-I, which allows websites to enumerate, manipulate, and access MIDI devices. Okay. The Magnetometer API, allowing websites to access data about the local magnetic field surrounding a user, as detected by the device's primary magnetometer sensor.

Leo: Well, there you go.

Steve: Which direction are you facing? And, for example, we've talked in the past about the battery API. And one of the problems was that it offered so much resolution, like I don't remember now, 16 bits of resolution, that if an ad could read your battery level, and then you as a user went to a different, completely unaffiliated website with exactly

the same 16-bit battery level - okay. 16 bits does not uniquely disambiguate you from the rest of the world. But when that is added to any other browser-sniffing, fingerprinting behavior that may be available, it represents a bunch of additional bits of - what would that be? Deentrification? There's a new word. Entropication, deentropication. Anyway.

Leo: There you go.

Steve: So magnetometer, which direction you're facing. Oh, and I just should mention that, as a consequence of that, it was decided that we did not need 16 bits of battery state, but it might still be nice to have some. So some of the browsers just fuzzed out the lower bits. They just said, we're not going to tell you exactly what the user's battery level is because you don't need to know. You only need to know if it's low enough that you shouldn't start a processor-intensive cryptocurrency mining on their browser right now or you'll burn up their battery.

Anyway, we also have the Web NFC API, which would allow websites to communicate with NFC tags through a device's NFC reader. That seems sort of useful. But not in Safari. The Device Memory API, allowing websites to receive the approximate amount of device memory in gigabytes. To me that seems like it's simply for the purpose of tracking someone. So, yeah, we could do without that one.

The Network Information API, providing information about the connection a device is using to communicate with a network and provides a means for scripts to be notified if the connection type changes. I would worry about that from a security standpoint. That sounds like it could be leaking something that bad guys might use. I mean, I'm sure they would have thought about that in the W3C Consortium, but still.

Oh, and speak of the devil, the Battery Status API, allowing websites to receive information about the battery status of the hosting device. I guess other browsers that have it are reducing its resolution. Apple says no, none of your business. And you know, it's interesting, too, because, I mean, there are enough of these that it means that - there are enough of these that Safari's not going to be supporting that it sort of limits websites' ability to depend upon those.

On the other hand, we've seen the numbers, and the Chromium browsers - which, by the way, are the browsers that support almost all of these, if not all of these, just because that's where Chromium is - they have the lion's share. So I guess, if nothing else, websites will have to be able to tolerate a lack of some of this functionality. Maybe they'll tell people, go get Chromium if you want to use these features. We also have Web Bluetooth Scanning. This seems like a bad idea. Allows websites to scan for nearby Bluetooth LE devices. What could possibly go wrong?

Leo: Yeah.

Steve: The Ambient Light Sensor. Web Consortium, really? Okay. Lets websites get the current light level or luminance of the ambient light around the hosting device via the device's native sensors. If any.

Leo: All of this is really designed to make the browser be like a built-in application.

Steve: Yes.

Leo: That's the whole point really; right?

Steve: Yes, yes. And we've talked about...

Leo: So you probably have many apps that do that. But you don't want your browser to do it.

Steve: Yes. And we've talked often about how the browser is becoming sort of the replacement desktop. Browser web apps are now really a thing. We also have, get this, the HDCP - remember, that's the high-definition copy protection, HDCP. We have the HDCP Policy Check extension which allows websites to check for HDCP policies, used in media streaming and playback. Okay.

The Proximity Sensor allows websites to retrieve data about the distance between a device and an object, as measured by a proximity sensor. Like, what? The user? Okay. The WebHID allows websites to retrieve information about locally connected Human Interface Device (HID) devices. And Leo, the Serial API, as in a serial interface, allows websites to write and read data from your COM ports, from serial interfaces, used by devices such as microcontrollers, 3D printers. So okay, you know, hook up your 3D printer. The website is able to talk to it directly and print something. But not in Safari.

The Web USB lets websites communicate with devices via USB. Interesting. The Geolocation Sensor, a more modern version of the older Geolocation API that lets websites access geolocation data. And maybe as a consequence of the still ongoing launch of the new third-generation GPS satellite that happened between your previous podcast and this one...

Leo: Yeah, that was cool.

Steve: ...that data got even more accurate. Or will, eventually, once it's up and in orbit. And then we have the User Idle Detection, lets websites know when a user is idle. So, okay. All of those things Apple feels are not worthy of implementation, they are mostly implemented in Chromium-based browsers. Firefox has a few. Chrome will have none. And of course Apple feels that, whereas third-party cookies were once the tracking mechanism of choice, as we know, they are readily blocked these days by savvy users. And much as I had my once naive hopes for the Do Not Track header - and Leo, you always chuckled at that, and you were right, as you could adopt the phrase DNT was DOA because it just never got off the ground.

Leo: No, no.

Steve: These days, in the place of cookies, the use of surreptitious fingerprinting has been steadily growing to the point that fingerprinting has now become the standard means of tracking users online for the advertising technology market. That's what the ad tech guys have ended up moving to is browser fingerprinting. So it's gone from sort of like an oh, gee, that's sort of an interesting idea, to no, that's the way we're going to track people now. And of course its growth has been driven by the browser vendors who have steadily been adding more or less subtle, sometimes less subtle, anti-tracking features to their browsers in the interest of enhancing the privacy of their users.

So anyway, Apple said that WebKit's first line of defense against fingerprinting is not to implement web features which increase fingerprintability in the first place, and offer no safe way to protect the user. And as for the traditional Web APIs which have been implemented in Safari for years, Apple has said it has been working to reduce their fingerprintability vector. Apple said that they had removed support for custom fonts, meaning that only presenting built-in fonts which are the same for all users with the same system. We know that enumerating installed fonts is one of the ways that websites were fingerprinting users.

They've also said they've removed minor software update information from the user agent string, so sort of fuzzing the version number so that it's not as specific. They've also removed the Do Not Track flag, which ironically was being used as a fingerprinting vector. If you had it turned on or turned off, it was generally sticky, and so it added one bit of fingerprinting disambiguation. So they just said we're going to turn it off for everybody, cannot be used. They've also removed support for any plugins on macOS. Other desktop ports may differ. And of course plugins were never a thing on iOS, so nothing lost there.

They require a user's permission for websites to access the Device Orientation and Motion APIs on mobile devices because the physical nature of motion sensors may allow for some additional device fingerprinting. And they finally said that they prevent fingerprinting of attached cameras and microphones through the Web Real-Time Communications API, you know, that's WebRTC. So anyway, there again is Apple being proactive. And it sort of goes along with the Apple user profile. You know, a little less focused on gee whiz features and a little more on just solid functionality and increase in privacy and security. So good for those guys.

Last Wednesday, during Apple's developer conference, Tommy Pauly, who's an Internet Technologies Engineer at Apple, explained that the fall releases of iOS 14 and macOS 11 would both be supporting the much beloved DNS-over-HTTPS and also DNS-over-TLS protocols. And of course that's good news for encrypted DNS and privacy, kind of. From what Tommy said, it doesn't sound as though it's going to be like an immediately consumer-friendly feature, at least not at the start. He said that developers could create apps to apply DoH and DoT settings for the entire operating system via network extension apps or MDM profiles, or to individual apps, or to an app's selected network requests.

So mostly security-aware users, sounds like they would have to add it as, like, an app to their device in order to get entire iOS device or Mac using the DoH provider of their choice. So I suppose somebody could create an app to enable that. But I don't get why not a setting in the OS's network settings. To me that would seem to be a far better solution. Maybe Apple regards it as like an advanced feature? I don't know.

Tommy said, quoting him, he said: "There are two ways in which encrypted DNS can be enabled. The first way is to use a single encrypted DNS server as the default resolver for all apps on the system." He says: "If you provide a public encrypted DNS server, you can now write a network extension app that configures the system to use your server. Or, if you use Mobile Device Management to configure enterprise settings on devices, you can push down a profile to configure encrypted DNS settings for your networks."

He says: "The second way to enable encrypted DNS is to opt-in directly from an app. If you want your app to use encrypted DNS, even if the rest of the system isn't yet, you can select a specific server to use for some or all of your app's connections." So, I mean, that's nice, but it's very developer-centric in the way that they're approaching it. So it's not as much as we would like, I mean, what you'd like is just to have your iOS device use Cloudflare or NextDNS or 1.1.1.1 or any of the increasing number of DoH servers. So

anyway, maybe they'll add that for us to iOS in the future. I guess we'll just have to wait and see how it goes.

And now for some less than wonderful news. I wrote here Mozilla plus Comcast plus DoH. I'm not sure whether that's strange bedfellows or strained bedfellows. But the disquieting news is that Comcast's Xfinity broadband Internet service will be joining Firefox's Trusted Recursive Resolver program, you know, TRR. And somewhat even more worrisome, it will be enabled, and Comcast will be the default DoH resolver for all Firefox users on Comcast's network. So in other words, for all Firefox...

Leo: What?

Steve: Yeah, I know.

Leo: That stinks. I'm sure they gave them a packet of money.

Steve: That was my first thought, Leo, was that Mozilla needs financial support. Whenever I buy something - I've never mentioned this. Whenever I buy something through PayPal, often there's a little "Would you like to toss in a dollar to Mozilla?" And I invariably say yeah, thank you. I'm happy to do that as a Firefox user. So I'm thinking, wait a minute. What has happened now is that, if you are a Firefox user, and you haven't done anything one way or the other, and you're at Comcast ISP, your Firefox will start using Comcast's DoH resolver by default.

Leo: That's not nice.

Steve: So let's enumerate the many things that are fundamentally wrong with this picture. First, recall that it was in 2014 that Comcast was caught injecting unsolicited advertisements into its customers' browsing. And the U.S.'s broadband privacy rules were killed by Congress in 2017 when Congress passed a bill allowing ISPs to collect and sell web surfing data.

Prior to President Trump's signing of that bill into law, as he immediately did with some flourish, the FCC's previous privacy rules would have required home Internet and mobile broadband providers to first obtain consumers' opt-in consent before selling or sharing web browsing history, app usage history, and other private information with advertisers and other companies. But lawmakers used their authority under the Congressional Review Act to pass a joint resolution ensuring that the rules "shall have no force or effect" and that the FCC cannot issue similar regulations in the future.

Then recall how last year, as we talked about at the time, Comcast went ballistic and patently lied in Congressional testimony about the move by the browser vendors to secure DNS. Remember that set of slides we reviewed at the time? It was like a PowerPoint presentation which was just full of nonsense that was created by Comcast and presented as part of Comcast's testimony to Congress.

And in a joint letter to Congress, of which Comcast was a signatory, the letter read: "We would like to bring to your attention" - you, Congress - "an issue that is of concern to all our organizations. Google is beginning to implement encrypted Domain Name System lookups into its Chrome browser and Android operating system through a new protocol

for wireline and wireless service, known as DNS-over-HTTPS (DoH). If not coordinated with others in the Internet ecosystem..."

Leo: So infuriating.

Steve: "...this could interfere on a mass scale with critical Internet functions..."

Leo: B.S. I'm sorry. Did I say that out loud?

Steve: Yes, good, I'm glad you did, "...as well as raise data competition issues." What? They said: "Google is unilaterally moving forward with centralizing encrypted domain name requests within Chrome and Android, rather than having DNS queries dispersed amongst hundreds of providers. When a consumer or enterprise uses Google's Android phones or Chrome web browser, Android or Chrome would make Google the encrypted DNS lookup provider by default, and most consumers would have limited practical knowledge or ability to detect or reject that choice."

Leo: Just as you, the members of Congress...

Steve: Uh-huh, but now they're doing exactly the same thing; right?

Leo: ...have no ability to detect B.S. at all. Oh, man.

Steve: "Because the majority of worldwide Internet traffic, both wired and wireless, runs through the Chrome browser or the Android operating system, Google could become the overwhelmingly predominant DNS lookup provider. Moreover, the centralized control of encrypted DNS threatens to harm consumers by interfering with a wide range of services provided by ISPs, both enterprise and public-facing, and others." Blah blah blah. Anyway, everyone gets the idea.

And so here's the thing that occurs to me when we talk about our ISP doing this. Remember that locally resolved DNS has never been a privacy issue as far as the wider Internet is concerned. When a consumer does use their own ISP's local DNS resolver, their queries never emerge onto the Internet. As we know, the local resolver is also a cache, which is able to satisfy most queries without asking anyone else. And when the local cache provided by the ISP doesn't contain the answer being sought, it and not the ISP's customer goes out onto the Internet to issue one or more queries of other DNS servers to obtain and then cache the answer on behalf of the customer.

In other words, the entire point of using DoH in the first place is specifically and only to blind the user's own ISP to the content of their DNS queries, which many people feel is none of the ISP's business. My point is there's no privacy implication as far as the rest of the Internet is concerned when a user leaves their DNS settings alone and uses the ISP's local cache. That doesn't go out on the Internet. It's that we're not happy about the ISP snooping on our DNS queries.

So yeah, we would rather send it to Cloudflare or to Google or to NextDNS or to whomever. We're telling our ISP to keep their nose out of our business. And based on past evidence and behavior, it's clear that our ISPs are every bit as desperate to know

our business as is every other entity out on the Internet. They want to track our behavior and our actions. So the speedy deployment of browser-based DoH, which frankly it's been surprising, I've talked about it in the past, just how quickly this happened. It was inevitable that ISPs would as quickly as they could be bringing up their own DoH servers to recapture those DNS lookups that they were rapidly being blinded to.

And what's most distressing is that we're going to witness the tyranny of the default, swinging Firefox users away from Cloudflare and NextDNS, which was what Firefox was going to be doing, and initially to Comcast and probably eventually to other ISPs, which are all going to want to bringing up their own DoH servers just for this purpose.

Mozilla's announcement said: "Mozilla, the maker of Firefox, and Comcast have announced Comcast as the first Internet Service Provider (ISP) to provide Firefox users with private and secure encrypted Domain Name System services through Mozilla's Trusted Recursive Resolver Program. Comcast has taken major steps to protect customer privacy as it works to evolve DNS resolution."

Eric Rescorla, Firefox's CTO, said: "Comcast has moved quickly to adopt DNS encryption technology, and we're excited to have them join the TRR program. Bringing ISPs into the TRR program helps us protect user privacy online without disrupting existing user experiences. We hope this sets a precedent for further cooperation between browsers and ISPs."

And for its part, Comcast has reaffirmed its commitment to the privacy of their users. Their explicit statement about privacy highlights three points. And maybe by saying it they're going to be bound by it. I don't know. They said: "As your Internet Service Provider, we do not track the websites you visit or apps you use through your broadband connection. Because we don't track that information, we don't use it to build a profile about you, and we have never sold that information to anyone. Two, we do not sell, and have never sold, information that identifies who you are to anyone. We also don't sell, and have never sold, your location data when you use our Xfinity Mobile service." And third and final: "We delete the DNS queries we have as an Internet Service Provider every 24 hours."

Leo: Well, that's good if they really do all that. I thought they were selling it. I mean, they have the right to sell it.

Steve: Yeah, they do have the right to sell it. They've said they wouldn't; but then we've all asked, okay, then, why are you fighting so hard to be able to?

Leo: Right.

Steve: So they are saying they're not doing it. So let's hope they aren't.

Leo: Well, you can still - it's just the default. So I can still choose Cloudflare or...

Steve: Yes, yes, yes, yes. And in fact that's exactly the point I wanted to make. I wanted to, for our Comcast-connected, Firefox-using listeners, be aware that, if you care, Firefox will at some point - and Mozilla hasn't said when, this is just an announcement of this agreement - it'll switch over to Comcast, away from whatever it was using before. So

you may want to override that and switch back to some non-ISP provider, if it's something that you care about.

Leo: Yeah. That's so weird that they would [crosstalk].

Steve: Leo, I know. Isn't it odd? Yeah. But it makes sense. I mean, ISPs don't want to be blinded. But on the other point, it was never really a security issue if you were using a local provider because your queries never left the ISP's network. They stayed in-house.

Leo: Right. Mr. G?

Steve: So we have the very welcome Facial Recognition and Biometric Technology Moratorium Act. It was proposed last Thursday by two senators and two representatives in Congress. Although a growing number of U.S. cities has already been taking action to ban the government use of that technology, with Boston just last week becoming the 10th U.S. city to do so, this bill, if it were passed, would put in place the first nationwide ban on facial recognition technology. And I don't know where this falls on political partisan lines, so maybe it's dead before it's even voted on. Who knows?

But the sense is that the use of facial recognition technology, we've been talking about it a lot on the podcast, has sort of sprung up unbidden because it had suddenly become feasible and practical. We got lots of computation. We got lots of connectivity. We got lots of storage in the cloud. So, hmm, what could we do with all of that? Oh, we've got lots of cameras, too. So let's hook everything together and, yeah, that sounds like a good idea. But the nation's lawmakers have never been given the opportunity to officially or unofficially weigh in on whatever limitations and guidelines, if any, should be imposed on this. It just kind of happened organically.

Leo: They probably want to know what Comcast thinks first before they vote. Have to wait for that letter. They'll tell you what to do.

Steve: We reserve the right. The newly proposed bill would "prohibit biometric surveillance by the federal government without explicit statutory authorization and to withhold certain federal public safety grants from state and local governments" - ah, there's the stick - "that engage in biometric surveillance." So that means that federal agencies would be barred from using biometric surveillance systems, which in addition to facial recognition can also include voice recognition. Additionally, it would prohibit the use of federal dollars to be spent on facial recognition technology and "condition federal grant funding to state and local entities, including law enforcement, on those entities enacting their own moratoria on the use of facial recognition and biometric technology."

So there's no time limit in place. It would continue until Congress passes a law to lift it, which is I think exactly the way it should be done. Just say no until we decide how we want to treat it, how we feel about it, what restrictions to put on it. And of course the privacy concerns around facial recognition have been one thing, and specifically, especially in today's climate, the inadvertent apparent racial bias repeatedly demonstrated by the technology, all of this came into focus earlier last week when the ACLU filed a complaint alleging that an African American man was arrested in Detroit after a facial recognition system falsely matched his photo with security footage of a shoplifter.

According to the ACLU, this is the first known example of misidentification through faulty facial recognition directly leading to a wrongful arrest. Apparently the technology was simply believed. The senior legislative counsel with the ACLU said: "No one should have to go through what the Williams family has gone through. It's past time Congress halted the use of face recognition and stopped federal money from being used to invest in invasive and discriminatory surveillance."

And coincidentally, a forthcoming research paper to be published by Springer Publishing of Berlin, Germany, bears the somewhat ominous title: "A Deep Neural Network Model to Predict Criminality Using Image Processing." What? Anyway, it describes AI algorithms that can predict crime based only on a person's face. I'm not kidding. In response, at last count, this was last night when I checked, 2,435 professors, researchers, practitioners, and students spanning all fields of anthropology, sociology, computer science, law, science and technology studies, information science, math, and more, including experts and academics from organizations including MIT, Microsoft, Harvard, and Google, had all signed an open letter denouncing this paper and calling it out for promoting racial bias.

And if anyone's interested their signed piece, I have a link. It was published on Medium.com last Tuesday entitled "Abolish the #TechToPrisonPipeline," with the subheading "Crime Prediction Technology Reproduces Injustices and Causes Real Harm." Anyway, it was published by the coalition - huh?

Leo: It's like phrenology. Like, oh, the bumps on this person's head tell me he must be a criminal type.

Steve: Exactly.

Leo: You can't look at somebody's face and say they're a criminal. Unless you just assume the skin tone tells all, which I bet is exactly what this is.

Steve: And that's the bias that we've seen, which is horrific.

Leo: Exactly what this is.

Steve: And as we previously noted on the podcast, Microsoft, Amazon, and IBM, which all have developed their own facial recognition platforms, have all recently banned the sale of the technology to police departments and pushed for federal laws to regulate the technology. So they're saying, yes, please regulate this. We have the technology. We've seen what it can do. We need some guidance here.

Senator Markley, who was one of the four who cosponsored this, said in his statement about the proposed legislation: "Facial recognition technology doesn't just pose a grave threat to our privacy, it physically endangers Black Americans and other minority populations in our country. At this moment, the only responsible thing to do is to prohibit government and law enforcement from using these surveillance mechanisms." So that's the happy proposed legislation. We will see how that fares through the U.S. legislative process. The not-so-happy legislative proposal is titled "The Lawful Access to Encrypted Data Act."

Leo: At least they're being honest.

Steve: And there it is, folks, yeah.

Leo: Right out in the open.

Steve: Yes, exactly, Leo. "The Lawful Access to Encrypted Data Act." Let's not call this the EARN IT, you know, we're going to take away Section 230 protection. No. So this new legislation was introduced by the U.S. Senate's Judiciary Committee Chairman Lindsey Graham, Senator Tom Cotton, and Marsha Blackburn. They argued that ending the use of - and they use the term "warrant-proof" - encrypted technology would "bolster national security interests" and "better protect communities across the country." They added that such encryption cloaks illicit behavior during criminal investigations into terrorists and other bad actors.

And of course, as we know, there's been no other single topic that has so captured my own interest and attention, along with the more recent question of managing the posting and presentation of counterfactual information on the web. The question of whether encryption should remain warrant-proof is the issue of the day. And it's going to be fascinating to watch this play out. As we know, we already have last year's EARN IT Act, also introduced by Senator Lindsey Graham, which in the show notes I called a "can of slime," whose language begins: "A bill to establish a national commission on online child exploitation prevention, and for other purposes."

Leo: It's the other purposes that scare me.

Steve: Yeah, and some other things. Don't worry about those.

Leo: Yeah, you don't need to - pay no attention.

Steve: Yeah, just think about the kids. And as we know, it threatens to revoke a non-complying company's legal protections under Section 230 which shields the company from criminal and civil liability for user-generated content. I much prefer, if nothing else, a straightforward heads-up fight over what we're really talking about.

Tom Cotton, one of the bill's three co-sponsors, said: "Tech companies' increasing reliance on encryption has turned their platforms into a new, lawless playground of criminal activity. Criminals from child predators to terrorists are taking full advantage. This bill will ensure law enforcement can access encrypted material with a warrant based on probable cause and help put an end to the Wild West of crime on the Internet."

So, okay, good luck with that. We all know, as has often been said on this podcast and others on the TWiT Network, that lawful users will obtain weaker privacy protection, while actual criminals who will be much more motivated to hide their elicit activity will simply switch to non-U.S. platforms that continue to provide warrant-proof encryption.

Oh, and in a somewhat bizarre carrot-and-stick, this latest attempt at legislation also directs - get this, Leo - our illustrious Attorney General Bill Barr to create a prize competition to award participants who create a lawful access solution in an encrypted environment.

Okay. So we know that Apple or Signal or Facebook are not going to be signing up for that prize, though I suppose this might be an effort to spur some innovation to get some smaller players to propose technological solutions since Congress has been unable to get any of the actual purveyors of this technology to even consider doing so. So anyway, we'll see how that goes. Who knows? None of this legislation ever seems to go beyond the proposal stage. And I assume that's because, presumably, those who query members in Congress about their feelings and probable votes never get even close to imagining that those bills have a snowflake's chance of actually passing.

So again, at least this is just, as you said, Leo, it's telling the truth. It's just saying, look, we want to end warrant-proof encryption. What say ye? And we'll find out because it's just been proposed.

Another legislative side, we have Michigan's state legislative house which just passed the - and I don't know why they had to do this, but I'm glad - the Microchip Protection Act. Even though forced RFID chip implants for workers is not yet, thank goodness, an issue, the state of Michigan decided to join seven other states in getting out in front of this one by formally and preemptively outlawing compulsory microchip implants for employees. Whoa. Big Brother, anyone?

The legislation notes that, although there are only a few known U.S.-based companies embedding microchips in their employees - huh? - several job providers could be following soon, they said, including some businesses in Michigan. So assuming that this legislation passes, which appears likely, Michigan would become the eighth state to explicitly outlaw compulsory chipping. The existing seven states are California, Maryland, New Hampshire, North Dakota, Oklahoma, Wisconsin, and Utah. All of those already prohibit the required implantation of a microchip in any person, for any reason, not just employees.

So I have a feeling that that's going to be, you know, taking all of this a step too far. So, I mean, I guess, like, "for any reason," that would even mean like convicts and criminals and things. It's like, no, we're not doing RFI tagging of people in this country. I imagine that this is just the beginning. We're a federation of states, so it's eight states now. I would imagine we'll see something nationwide before long.

Oh, and I wanted to mention, just sort of as a heads-up, not to forget about VirusTotal. A recent blog posting, a VirusTotal blog posting, announced a new addition to Google's multi-malware scanner. I got a kick out of the title. It was kind of some clever programmer-ese. The blog was titled "VirusTotal += Cynet." For those not conversation in coding, the "+=" expression happens to be one of my favorites. It's a simplification of "a=a+b." It should always be the goal of computer languages to both minimize writing errors and maximize reading ease. So shortening "a=a+b" to "a+=b" is vastly easier to both write and to read. It's easier to read because it explicitly reveals the programmer's intent to increase "a" by "b."

Now, it's true that the longer version does the same. But the intent is much less clear. And especially if the "a" expression happened to be some really complex thing. You don't want to have to type that twice on both sides of the equal sign and then add a "b" to the end. And if it's really long and complex, someone reading the code looks at it and thinks, wow, what's going on? And then they have to very carefully make sure that the complex expression is in fact the same on both side of the equal sign. Anyway, "+=" is a win.

So Cynet is a new addition to VirusTotal, and the VirusTotal blog quoted Cynet, saying: "Cynet 360 is an autonomous breach protection platform that includes multilayered antimalware capabilities including AI-based static analysis, process behavior monitoring, memory monitoring, sandboxing, and granular whitelisting, interlocking all together to

protect against malicious executables, exploits, scripts, macros, malicious process injection, and other fileless attacks."

They said: "Cynet 360 protection ranges across the entire malware lifecycle, identifying malicious attributes," blah blah blah. Anyway, so that's just an ad for them. But they are a good standup company, and they will now be contributing their input into things dropped into VirusTotal. I wanted to mention it mostly to remind our listeners about VirusTotal, a service which I use and, frankly, would hate to be without.

Just the other day I needed some piece of long-abandoned software. I don't now remember what it was. So I went digging around the 'Net, and I ostensibly found it on Vetusware, V-E-T-U-S ware, which is Vetusware.com, which is an abandonware site that tends to collect abandonware, you know, arguably software which was once commercial, but which has been long abandoned by its publisher. And I think there's actually been some legislation protecting those who use abandoned software, saying yeah, you have a legal right to do so. Much as I wanted to trust the download that I found...

Leo: Looks like the site has been abandoned.

Steve: Yeah. Well, now, Leo, it looks a lot like mine, so easy there.

Leo: I love the use of ASCII text characters for borders and shadows. That is a pure ASCII [crosstalk], yes.

Steve: Yes, I do, too. Anyway, so much as I wanted to trust the download, this is precisely how malware gets into people's machines, you know, downloading something that you think is okay, and it's not. So I downloaded whatever it was I needed, inspected the package, checked file signatures, dates, and so forth. I did as much due diligence as I could. And for me, a key part of any such checking is always to drop any questionable file onto VirusTotal over at VirusTotal.com. Sometimes it instantly knows it because it does a hash of it, and it had already checked the file of that hash. But it's sometimes a good idea to refresh that.

And in this case I did, and I watched 73 different virus, you know, AV antimalware engines all reinspect that thing and say, yeah, far as we know, nothing wrong with this. And as it turns out it was useful to me, and it was benign. So anyway, I'm not suggesting that anyone should regard anything VirusTotal tells you as gospel. But it's one additional piece of useful telemetry about anything you might have reason to question. So just remember that it's always there, and it's free to use, and just a great resource for security-savvy people on the web.

A little bit of miscellany regarding Edge on Win7. I mentioned it last week, and I heard you confirm it the day after, Leo, with Paul and Mary Jo. As it happens, since then and now, I updated a Windows 7 machine of mine which had been offline for quite a while, and it did receive an offer to install Edge. So Windows 7, we'd like to install Edge. The Win update, it was interesting the way it came in. It came in under Windows Update, which showed me that two important updates were available. So it was important, but Edge was not selected for me by default. So if I didn't do anything, if I just did the automatic express updates, I would not have received it. But if I looked at the itemization of updates, saw that it was important, I would have gone, oh, yeah, I'd like to have Edge on my Windows 7 machine. I could have selected it and installed it. I did that, and it installed it, and it worked.

What I did think was interesting was that it was rather aggressive about copying Chrome's settings. Since it of course is a Chromium-based browser, it really wanted to suck everything out of Chrome and take over for it. And I did allow it to do so because I was just sort of curious to see how it would do. And sure enough, it installed all the same add-ons, including Last Pass and uBlock Origin that I both had installed as add-ons in Chrome. So anyway, I know that there are still lots of Windows 7 users out there, and I don't know why somebody who is using Chrome would also want Edge. It's just one more browser. But, you know, Microsoft has been telegraphing that they're going to be adding a bunch of features, like vertical tabs. So maybe.

I also wanted to note, for those who are old-school NNTP newsgroup users, apparently there are about a thousand of you, since I think about a thousand copies of the Gravity newsreader have been downloaded since I added it to GRC's list of downloads. I wanted to just give a quick heads-up that there's a new release of that. I fixed a bug in the one that we had online, 3.0.9. The bug was not mine. A couple of weeks ago Gravity users, including myself, were unable to open several recent postings in the grc.spinrite.dev newsgroup, where I'm spending all my time. Gravity would just crash. It would just close, bang, whenever you touched on one or two of those posts.

And this is really where community-driven open source software really comes into its own. Gravity, as we know, was abandoned, I think in 2010, quite a while ago. But it is by far my favorite way of interacting with newsgroups. And newsgroups are, by far, the best way I've ever found of working with a community of people who are interested in participating in the development testing of new code, which is what we are very busy doing at the moment.

So when my trusty old Gravity died a couple weeks ago, well, actually first remember that it died when the year switched to 2020. There was just some code in there that, you know, it wasn't Y2K, it was Y2K20. It just died. So I needed to open up and fix it at that time. I learned a little bit about Gravity, brought it into the fold. I now have its source tree living with me. And so when it crashed two weeks ago, I went back in, and I found a longstanding problem.

Somebody, the original author, I mean, this thing had never worked. If you embedded two Base64 encodings into a posting, which is the way Thunderbird does things because Thunderbird has a newsreader also, and you tried to open that in Gravity, it would crash. The original programmer forgot to put the length of a string into one of the string abstractions. You know, when have we ever heard of that happening before? And sure enough, it would crash. If the Base64 decoder ran off the end of the first one and into the beginning of a second one, it would hit an equal sign which it regarded as a reserved character because Base64 uses that for padding, and it would explode. So anyway, I fixed it. And for anybody who did download a copy of Gravity, you'll want to go update your copy just so that you have the latest and greatest.

Also, a bit of feedback from a listener, BlueLED. He tweeted me: "@SGgrc Can you tell us again what was the brand and model of the switch/router you were using for traffic control? It was not a common brand, as I remember."

Okay. So the software is pfSense, P-F-S-E-N-S-E. It is open source, based upon FreeBSD Unix, and it is wonderful. It can be loaded onto any Intel-based PC hardware which FreeBSD supports. And it supports everything, and going way back in time. So you do need to have some fast network interfaces if you're going to make it your router because of course all of your LAN traffic would be coming and going through it. So that's one requirement. But it is perfect for a DIY user.

You could also get any cute little inexpensive Intel-based fanless turnkey box and load pfSense onto it, and you're up and running with a feature-packed router. But for

someone who wants a pfSense-based appliance, the company Netgate offers a turnkey, ready-to-go pfSense-based router. Actually, they have a family of them. The smallest, least expensive is the SG-1100, no relation. It sells for $179. It's the one that I'm using at several of my locations. I think it runs up to 895Mb, so just shy of a gig of measured throughput. So it should handle anybody who has a 300MB connection as I do from Cox. Anyway, Netgate, N-E-T-G-A-T-E, dot com. And I've got the links in the show notes for anyone who's interested.

Work on SpinRite's AHCI driver is continuing. We're making rapid progress. I don't have anything specific yet to announce, though I did notice that more than 1,500 people had grabbed a copy of our new InitDisk utility, which I mentioned last week, for reformatting USB drives. So I'm glad it's been useful to people.

Leo: Steve?

Steve: So Ripple20 Too, as in also, more, mucho.

Leo: Oh, boy. Yikes.

Steve: Details, yeah. Last week was our disclosure and discussion of the very worrisome Ripple20 discovery that a highly used for several decades embedded TCP/IP stack was riddled with at least 19 vulnerabilities, several of the worst being very serious no-user-action-required remote code execution enablers. I entered that discussion last week by noting that the Internet's already target-rich environment just got a whole lot richer. So we're back to this important topic this week as a great many more details of the impact of this sweeping, industry-wide problem are coming to light. And the sense is we still don't know the full scope of it.

But our often-cited friends over at BleepingComputer have pulled together the most comprehensive coverage I've seen. So I'm going to start by paraphrasing a bit of what they've written because it nicely characterizes the industry's present situation. Bleeping Computer wrote: "The dust is far from settled following the disclosure of the 19 vulnerabilities in the TCP/IP stack from Treck, collectively referred to as Ripple20, which could help attackers take full control over vulnerable devices on the network.

"Treck's code is fundamental for the embedded devices it is implemented on because it bestows network communication to them and is present on gadgets used in a variety of sectors: technology, medical, construction, mining, printing, energy, software, industrial control systems, telecom, retail, and commerce.

"The company has notified its customers and issued patches, but a week after the Ripple20 announcement from security research group JSOF, the full impact still remains unclear. This is because Treck's code is licensed and distributed under different names, or serves as a foundation for a new network stack."

And I'll interject that it turns out that it's also internationalized, resold, and rebranded under completely different names, as well. So it's very possible that a mid-chain OEM might not even know that their system is using vulnerable code that originated with Treck. I mean, this is that whole - the reason they called it "Ripple" is that the nature of today's supply chain, where somebody could get this embedded in a chip, somebody then uses the chip, oh, look, it's IP-enabled. It's like, yeah, they don't even know that it's Treck code in there because that affiliation didn't survive the embedding process.

So, "Concerted efforts from national-level cybersecurity agencies and private companies in the field are ongoing to identify businesses with products vulnerable to issues in the Ripple20 vulnerability set. What is clear at the moment," they wrote, "though, is that the healthcare industry is particularly affected and should be on high alert." And of course this is interesting because it's sort of the nature of embeddedness. There is some provider that created the core embedded chips which either one or multiple major biomedical firm used. And, you know, an engineer chose it. It did the job. It was exactly that. It was a turnkey solution that would connectivity-enable their infusion pumps, for example. And then they just kept using it.

And they used it for years and years because it stayed available, and it worked, and they never had any problems with it. And so as they evolved their line of biomedical equipment, they just kept using it. And when they added some other piece of equipment, they said, hey, let's just use the same chip. Why not? It works. And the Treck products do work. The problem is they've also got vulnerabilities.

So the healthcare industry has been put on high alert about any of their connected stuff. Elad Luz, head of research at CyberMDX, which is a company focused on security and medical devices and is involved in identifying vulnerable products, told BleepingComputer that their initial investigation placed the healthcare industry's exposure at more than seven times that of manufacturing. And that was confirmed with some numbers by the security firm Forescout, who is also involved now in the effort. On the day of the disclosure, Forescout revealed that there were six times more vulnerable healthcare-related equipment than in the retail sector.

So what are the numbers? They have identified 52 - I have a hard time saying it - 52,935 devices matching the Treck signatures in the healthcare vertical segment. 8,347 devices in the retail segment. 7,333 devices in manufacturing. 5,904 devices in government. 5,225 in financial services. And generically an additional 11,346 in other fields. They explained from their analysis, Forescout did, that the most common device types running Treck stacks were infusion pumps, and then second were printers, then UPS systems, networking equipment, point-of-sale devices, IP cameras, video conferencing systems, building automation devices, and industrial control systems.

They added that to exploit Ripple20 vulnerabilities, an attacker needs a direct connection to an affected device or a routed path to one on an internal network. This means devices directly connected to the Internet, of course, are those most at risk, just like routers. An attacker could target these devices first, compromise the device, then move laterally within the network to access or infect other devices. So, for example, the last thing you want is your IP camera to be on the Internet with a port. Even if the port is not a server. If it's not visible, it is still, if it's got a port connected, these vulnerabilities can get there because, for example, an unsolicited DNS reply can be used to compromise one of these Treck stack-based devices.

So as an example of the impact of the vulnerabilities, Forescout did a series of Shodan searches for 37 specific known vulnerable device models where the model number was known, encompassing 18 different vendors. That included printers, IP cameras, video conferencing systems, networking equipment, and industrial control devices. That Shodan search revealed 15,000 currently Internet-connected instances of these devices with Treck stacks that could be immediately compromised by anyone on the Internet. And when you include networking equipment and industrial control systems, especially when you hear what some of these industrial control systems are, that's a problem.

So compared to the number of known vulnerable devices, estimated in the many hundreds of millions, 15,000 is not a big number. But that was only a scan for specific known vulnerable signatures. And after all, it only takes one exposed device on the network of, say, a nuclear power plant, or managing a nation's electric power grid, to

ruin everyone's day. It is estimated that the actual number of publicly exposed Internet-connected devices is a great deal higher. They just aren't readily identifiable from a Shodan scan, or they have not been identified yet. So that remains for the short-term foreseeable future.

BleepingComputer assembled a nice, almost alphabetical list - I alphabetized it a bit - of known affected manufacturers and their devices. To give our listeners some sense, we have the big iron provider, Aruba Networks. A preliminary advisory from them based on an initial investigation is available from them. It lists their Level 2 and Level 3 switches produced under the Aruba or HP ProCurve brand names.

Baxter U.S. They're a healthcare company. They announced that some of their Spectrum Infusion System's Wireless Battery Modules are impacted by Ripple20 because they run something known as the Digi NET+OS which uses Treck's TCP/IP stack. And we'll talk about Digi in a second. They list five different WiFi wireless battery modules with b, b/g, a/b/g/n and so forth. And we know what those are. Those are the WiFi designations. So they've got the Ripple20 Treck stack in them. And I don't know what mischief somebody could get up to, but it's not a vulnerability you probably want in something which is presumably infusing drugs into you at a prescribed pace.

Leo: Yeah. Let's not hack that.

Steve: Yeah. Braun, another medical and pharmaceutical device company, notified their purchasers that vulnerable Treck code is present in their Outlook 400ES Safety Infusion Pump, and no other products are affected in the Ripple20 issues. A Braun advisory stated that: "To date, Braun has received 24 patches from Treck to resolve vulnerabilities in the software." They said: "We've analyzed the patches and determined that 20 of them are not applicable to the Outlook 400ES platform. Four remaining patches continue to be analyzed to determine the scope." So that's good. That says that Braun purchased directly from Treck. Treck has been responsible. They've resolved the problem, and they're notifying their direct customers. Of course, the problem is with the indirect customers that Treck has no way of notifying. And we don't know if the direct customers have a way of notifying their down-the-chain OEMs.

Beck/HMS Industrial Networks is another. And rather than going through this in detail, to give our listeners a sense, Carestream has a bunch of products that are vulnerable. Caterpillar is not being forthcoming with details, but an undisclosed number of Caterpillar's products are known to be vulnerable. I don't know what that means. Cisco also. And, boy, Cisco is apparently a user of this Treck stack. Routing and switching, enterprise and service provider. ASR 5000 series routers. GGSN Gateway. GPRS Support. IP Services Gateway. These are just a whole bunch of Cisco. System Architecture Evolution Gateway. So yikes. And they're still investigating and will be showing more of them.

Dell said Dell products are inherited from an Intel component that is present in Dell Client Platforms and from Teradici (T-E-R-A-D-I-C-I) firmware and remote workstation cards in Dell Precision and Dell Wyse (W-Y-S-E) Zero Client products. So, interesting that they're saying they inherited it through an Intel component. That's interesting. They've released fixes and encourage customers to prioritize updating their systems to the latest firmware.

Also Digi International has something called NET+OS 7, which is a resold platform. It's got problems, and there's one, two, three, four, five, six, seven, eight, I don't know, 12 or so of those, and a whole bunch of model numbers on each of those. New firmware versions are available. They're being responsible. But it's up to people to install them. And the multinational power management company Eaton announced in a security

bulletin that its products rely on Treck's library to implement IPv4, IPv6, UDP, DNS, DHCP, TCP, ICMPv4. Remember that all of those were vulnerable protocols of Treck's library, and they are vulnerable to multiple Ripple20 issues.

Unfortunately, these things have sort of scary names. The CL-7 voltage regulator control. The Form 4D recloser control. The Form 6 recloser control. Edison Idea and IdeaPLUS relays, all variants are vulnerable. The metered input power distribution units, the metered outlet power distribution units, the managed power distribution units, and the high-density power distribution units. So, yeah. Let's hope there's no grid outages anytime in the near future.

Green Hills Software. HCL Technologies. Hewlett Packard Enterprise. HP Inc. and Samsung. HP Laser, LaserJet Pro, the HP Neverstop Laser. Maybe it should stop while you update its firmware. The Samsung ProXpress, the MultiXpress, the DeskJet, OfficeJet, OfficeJet Pro, Ink Tank, Smart Tank. You're in the tank. No. Anyway. You definitely want to see about getting updated firmware. We've talked about the problem with printer firmware vulnerabilities in the past. It is a favorite pivot point. There was just also something recently about printers opening a port to the public Internet. You don't want it to be one of these printers, or one that hasn't had its firmware updated, because that would be bad. That's a way to get in and pivot into your network.

Also Intel. Some versions of Intel's Active Management Technology (AMT), if I'm not mistaken, that's part of the baseband technology in the motherboard. So which is to say that Intel is using the Treck stack in the motherboard baseband firmware to do all of the active management technology which, as we know, cannot be turned off. So let's hope the people have those servers, the ports that have those servers are not facing the public on the Internet. That would not be good.

MaxLinear, Rockwell Automation, Schneider Electric. Lots of vulnerabilities across Schneider's line of UPS devices. Teradici I mentioned before. They've acknowledged the problem, and I'm sure they're notifying their customers. Xerox B205, B210, and B215, whatever those are. They've got firmware updates, but of course you've got to go get them. And then Zuken Elmic, a company that distributes Treck's stack under the name KASAGO (K-A-S-A-G-O). So if you don't know Treck, you may know that name. That's an example of one of these redistributors of the Treck stack under a different name.

So anyway, as we said last week, this is the vulnerability that I have a feeling we're going to be hearing about. I'm not going to spend a lot of time talking about it again. The problem is the bad guys who find it are going to be quiet about it. They're going to be using it as a way of getting into people's networks and taking advantage of what is unfortunately a longstanding, widely used stack in embedded IoT devices. It's definitely something bad guys are going to be probing for. So the advice for our listeners is this is just a good time to think about all the stuff you've got - printers, routers, IP cameras. Just go be proactive. Too few of these things update themselves, especially if they've been around for a long time. If there's new firmware for them, it's worth installing.

Leo: Now, did I miss it, or did you skip the Picture of the Week?

Steve: I skipped it. I didn't think it was that noteworthy.

Leo: Oh, okay.

Steve: It's just kind of fun.

Leo: Well, you know, it's there. I'll just show it for those of you watching at home.

Steve: Yeah.

Leo: People were saying in the chatroom, oh, you know. Then I thought, no, I don't remember that part. And since I'm the guy who's pushing the buttons, I should remember it. It's just a picture of a hard drive. Guy's looking at it with a magnifying glass, writing the ones and zeroes. Silly. Silly. At least he's wearing gloves. I think he's wearing gloves. I think he's wearing gloves.

Steve: It probably was - it was a photo that someone tweeted me. I thought it was a fun Picture of the Week. The person is wearing greaseproof gloves, and they're probably dustproof. So presumably it was meant to be a cleanroom environment where they'd opened the drive. And of course you can't actually use a microscope. Now, there were actually, in the old days, in the MFM drive days, there was a fluid you could put on the surface of a disk.

Leo: No.

Steve: Yes. And it would show - it would resolve the magnetic poles of the data.

Leo: So it was magnetic fluid or something. Holy cow.

Steve: You actually could, with a magnifying glass, you actually could read it. The bits were that big.

Leo: That's wild. That's amazing.

Steve: But that stopped being true a long time ago.

Leo: Long time ago. Steve Gibson, he will always be big in our imagination and our esteem. Every week we do the Security Now! show. You can tune in and watch us do it live around about 1:30 pm Pacific on Tuesdays, 4:30 Eastern, 20:30 UTC. So Tuesday afternoon or evening, tune in by YouTube, well, it's on YouTube Live, but actually the best thing to do is go to our website, TWiT.tv/live, and you can pick your stream. No more Mixer, I'm sorry to say, but there's still plenty of other streams there. TWiT.tv/live.

Steve's got a number of different unique versions of the show. His world-famous 16Kb edit, which just sounds wonderful. And then he also has the transcripts, which no one else has. That's so you can, you know, he has it all transcribed by Elaine Farris so as you can read along as you listen. And he has 64Kb versions. GRC.com. While you're there, pick up SpinRite, well on its way, I might add, well on its way to 6.1.

Steve: It's underway.

Leo: If you buy it now you'll get the next edition for free. You'll also get to participate in the beta testing. That's SpinRite at GRC.com. He's got lots of other free stuff there like ShieldsUP! and so forth. So go on over and see that. We have the show, audio and video of the show at our site, TWiT.tv/sn. It's also on YouTube in a store-and-forward fashion. Best thing to do, though, use that fancy RSS Reader you've got there that's known as a podcast program, a podcatcher. Subscribe to Security Now!, and it'll automatically download it the minute it's available of a Tuesday evening, so you'll have it for your Wednesday morning drive down to the store to pick up more brisket.

My brisket's frozen. For the whole show it's been at 137 degrees. Stalled, they call it. It's stalled. You know, it's interesting. There's a whole science about that. People weren't sure why brisket would stall for a couple of hours, just stop at a temperature. Turns out it's evaporating the liquid. It's actually cooling itself so it can't get any hotter. Once all the liquid's gone - I told Lisa this, she believed me - it will start to heat up. So I'm going to go home and watch my brisket. So everybody, thank you for being here. Thank you, Steve. Have a wonderful week. We'll see you next time on Security Now!.

Steve: Thanks, buddy.

Leo: Come on over for some brisket. Come on over. It'll be done in about 18 hours.