SECURITY NOW!

Transcript of Episode #765

An Authoritarian Internet?

Description: This week we add Bruce Schneier's thoughts about the theoretical feasibility of contact tracing apps. We touch on our government's feelings about DNS over HTTPS. We look at yet another wacky way of exfiltrating data from an air-gapped computer. We examine a new vulnerability that has already damaged some large high-profile enterprise infrastructures. We note Adobe's latest round of critical updates, another welcome service coming from Mozilla, a dispiriting bit of over-the-top political correctness from the U.K., and Google's plans to clean up the mess which is the Chrome Web Store. We share a bit of errata, miscellany, and SpinRite news, then take a look at China's proposed changes to the fundamental operation of our global Internet.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-765.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-765-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We'll talk about RDP scanning. It's way up. Bruce Schneier talks about why contact tracing apps are futile. And we'll talk a little bit about SaltStack and a big security flaw for a lot of companies out there. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 765, recorded Tuesday, May 5th, 2020: An Authoritarian Internet?

It's time for Security Now!, the show where we cover your privacy, your security, your safety online with this guy right here, Steve Gibson of GRC.com.

Steve Gibson: Hi, Mom.

Leo: His mom? No. Mom's watching from up there. You might be saying hi to my mom. I don't know if she watches this show.

Steve: Oh, I hope not. No, that would...

Leo: It's a little geeky.

Steve: It would do some damage. So as promised, we're going to talk about China's proposal for how to fix the Internet.

Leo: Yikes.

Steve: And thus as a consequence of some of the features of it - which maybe it's a little reactionary, but we'll let our listeners judge - I titled this "An Authoritarian Internet?"

Leo: Oh, boy. Yeah.

Steve: But it's got some interesting features. But yeah, it's like, well, not going to happen. But it is May 5th as we're recording this, Episode 765. We've got a bunch of other stuff to talk about that I think will be interesting. We're going to add Bruce Schneier's thoughts about the theoretical feasibility of contact tracing apps. We touch on our own government's feelings about DNS over HTTPS. We look at yet another wacky way - yes, another one - of exfiltrating data from an air-gapped computer, brought to us by the same guys who keep doing these things.

We examine a new vulnerability that has already damaged some large high-profile enterprise cloud infrastructures. We note Adobe's latest round of critical updates that they dropped last Tuesday. The news of another welcome service coming from Mozilla. A dispiriting bit of over-the-top political correctness from the U.K. And Google's plans to clean up the mess which is the Chrome Web Store. We then share a bit of errata, some miscellany, and as I thought I was going to be able to promise, some SpinRite news. Actually very welcome SpinRite news.

Leo: Oh, that's exciting, yeah.

Steve: Yeah. It was a good week. And then we're going to take a look, as I said, at China's proposed changes to the fundamental operation of our global Internet. So Security Now! 765 coming up.

Leo: As usual, great big packet of informational stuff. Steve?

Steve: So I didn't mention before that we have a little bit of a fun Picture of the Week. This actually appeared. This was a story in Spectrum.IEEE.org Tech Talk column on computing. The headline of the story, dated April 10th: "Cobol Programmers Answer the Call to Shore Up Unemployment Benefits Systems." And the sub is: "Retirees and newcomers want to help fix old software overloaded by new claims caused by the coronavirus pandemic."

So of course the problem is that some of these systems are, well, written in COBOL, and the population of people who are still fluent - I guess it's still actually around. But anyway, they need more than they have. So we've got old-timers coming out of retirement saying, yeah, I remember how to read that source code. And they're coming back to I guess fix things that are broken and...

Leo: That's hysterical. I love it. I love it.

Steve: ...crashing under the strain. I just got a kick out of that. Yes, we still need you, old-timers. Somebody someone will need me to fix some assembly code somewhere, so I'll go, oh, yeah.

So Bruce Schneier decided to weigh in. Of course we all know Bruce. He's a well-known cryptography and privacy guy. He blogged, and he begins his blog posting, which was titled "Me on COVID-19 Contract Tracing Apps," by writing, he starts: "I was quoted in BuzzFeed." And then he quotes himself being quoted. BuzzFeed said, quoting him: "'My problem with contact tracing apps is that they have absolutely no value,' Bruce Schneier, a privacy expert and fellow at the Berkman Klein Center for Internet & Society at Harvard University, told BuzzFeed News." Again quoting: "I'm not even talking about the privacy concerns; I mean the efficacy. Does anybody think this will do something useful?" He says: "This is just something governments want to do for the hell of it. To me, it's just techies doing techie things because they don't know what else to do."

Okay. That was the quote from BuzzFeed. Then he elaborated in his blog posting. He says: "I haven't blogged about this because I thought it was obvious. But from the tweets and emails I have received, it seems not. This is a classic identification problem, and efficacy depends on two things: false positives and false negatives."

So then he has two bullet points. False positives first. "False positives: Any app will have a precise definition of a contact. Let's say it's less than six feet for more than 10 minutes. The false positive rate is the percentage of contacts that don't result in transmissions. This will be because of several reasons. One, the app's location and proximity systems, based on GPS and Bluetooth, just aren't accurate enough to capture every contact. Two, the app won't be aware of any extenuating circumstances, like walls or partitions. And three, not every contact results in transmission. The disease has some transmission rate that's less than 100%." And he says in parens: "(And I don't know what that is)."

Then, for false negatives: "This is the rate the app fails to register a contact when an infection occurs. This will be because of several reasons. One, errors in the app's location and proximity systems." Which of course was the same as the first one. "Two, transmissions that occur from people that don't have the app." He says: "Even Singapore didn't get above a 20% adoption rate for their app." He says: "And three, not every transmission is a result of that precisely defined contact." He says: "The virus sometimes travels further." And of course we also know that it also can be transmitted through surface contact, not person-to-person.

Anyway, he said: "Assume you take the app out grocery shopping with you, and it subsequently alerts you of a contact. What should you do? It's not accurate enough for you to quarantine yourself for two weeks. And without ubiquitous, cheap, fast, and accurate testing, you can't confirm the app's diagnosis. So the alert," he says, "is useless. Similarly, you take the app out grocery shopping, and it doesn't alert you of any contact. Are you in the clear? No. You actually have no idea if you've been infected. The end result is an app that doesn't work. People will post their bad experiences on social media, and people will read those posts and realize that the app is not to be trusted. That loss of trust is even worse than having no app at all."

He says: "It has nothing to do with privacy concerns. The idea that contact tracing can be done with an app, and not human health professionals, is just plain dumb." Okay. And of course that's Bruce Schneier, who's Mr. Tech and crypto and so forth. So that's how he ends his blog.

Leo: And we've talked about this. It doesn't have to be perfect. We're trying to just bring the R-naught, the contagion rate below one; right?

Steve: Correct.

Leo: So even if it was 50% reliable, it would have value; wouldn't it? Or maybe I don't understand it.

Steve: Well, so for me the numbers out of New York are really fascinating because they're still in, I mean, they have been from the beginning in as stringent a lockdown as possible with Andrew Cuomo, the governor of New York State, spending an hour every day sort of as the father confessor, I mean, just like really urging people to stay at home, stay away from other people. And with that they've only managed to get it barely below 1.0, which to me is fascinating. I mean, just independent of what we would think. In part of the state it's 0.8, they've calculated. And I think in the north it's 0.9. But, I mean, they're doing everything they can. And the epidemiologists see that, and that's why what we're now hearing is everyone who's informed expects we're going to see a significant increase in cases about two weeks from now because this thing just wants to get transmitted. I mean, it is trying to, hard.

Anyway, there was an interesting thing that he also said that I thought was interesting. At the end of last week, Cuomo explained that the previous day - he was talking about their big project to ramp up human contact tracing; right? So here's Bruce saying forget about the app, you need trained medical professionals. So Cuomo is going to do this. Bloomberg is stepping up, and he's organizing the project somehow.

So he said at the end of last week, he said that the previous day, which was last Thursday, the state identified 4,681 new, brand new cases of coronavirus, and that about the same would be happening again that same day, and pretty much every day. He explained, I mean, and I was amazed he could do this with a straight face, he explained that effective contact tracing required that every one of those 4,681 people be interviewed to determine the identities of everyone they had come into physical proximity to over the past 14 days. And then...

Leo: They're going to telling people you've got to keep a journal of everywhere you go and what you do and who you see and who you meet for two weeks.

Steve: Well, and not just if you're infected, but everybody, because you never know...

Leo: You don't know.

Steve: ...when you're going to suddenly be tested positive.

Leo: Right.

Steve: And then all of those people, all of the people who the 4,681 people who tested positive on one day came in contact with for the previous 14 days, all of them need to be interviewed and perhaps placed into isolation, presumably interviewed and tested, and perhaps placed into isolation. 4,681 per day. And then that expands, as I said, to all the people that they came in contact with. And that has to happen every single day. Does anyone believe...

Leo: I heard Mike DeWine, the governor of Ohio, yesterday say, "Hey, we got 1,800 contact tracers. We're in great shape." I don't think that's going to do it.

Steve: I mean, no. And in fact the most recent estimate is that we will need 300,000 human contact tracers. So that'll be good for unemployment, but you'd have to train them up. But the point is, Leo, 4,681 people in one state every single day then interviewed to determine everyone they may have had contact with for the previous two weeks, and then interview all of them. It's not possible. And so that's my point is that this thing has escaped. It is out there. And my feeling is it doesn't matter at all when we so call "reopen." It's how we reopen that matters. And from a standpoint of the damage we're doing to the economy, the sooner the better.

But as we were just talking, I guess it was before we began recording, about how you're out now in Petaluma, and people are not wearing face protection. They're not covering their faces. And so I'm of the opinion that we absolutely need to continue to be careful, but that you do as much commerce as you can while being as careful as you can be because this thing appears to be insistent upon finding people to infect. If with New York doing everything they possibly can, they've barely managed to sneak it below an R-naught of one, then the moment they relax restrictions it will go above one. And then we risk having this thing get away from us again. I mean, it's a challenge.

So anyway, I thought it would be interesting to share Bruce's sentiments about software. My feeling is it'll be fun to play with. I'll install it, and I'm not worried about privacy concerns. And if it goes off and sounds the alarm, it's like, oh, okay. And the problem is it doesn't tell you who you got it from. It just tells you sometime in the last two weeks somebody who tested positive was near enough to you that you exchanged Bluetooth tokens, and good luck to you.

Leo: Yeah. Well, this is why you want scientists and epidemiologists and physicians to make these decisions, not politicians or, with all due respect, podcasters. This is crazy and tough, and I don't know what's going to happen. I really don't.

Steve: Yeah. I don't think I've posted the link. There's a site that does a really good running three-day average. As I said a long time ago, when I see people talk about the number of cases they've had, it's like, no, it's the number of cases you've identified. And that's why, you know, very early on I said unfortunately, and ghoulish as this is, death is the only real count we have. And I've been worried about India because they are late to the game, but they have a very dense population. And anyway, the site that I've been following, I've looked at India, and it's just beginning to go exponential. So the good news is they're late, so they've had the benefit of learning from the rest of the world that achieved critical mass of this virus months before. So hopefully they will be able to keep it under control. So interesting times we're in.

Leo: No kidding.

Steve: Yeah. So I said here "DHS's CISA says no to third-party DoH."

Leo: All there.

Steve: So we have abbreviation soup. DHS is of course the U.S. Department of Homeland Security. CISA is the Cybersecurity & Infrastructure Security Agency, which is an agency within the DHS. And of course we all know that DoH is DNS over HTTPS.

In a recent four-page memorandum, the U.S. Department of Homeland Security reminded all federal CISOs, Chief Information Security Officers, that they must not use the DoH services which are coming to all our web browsers. I have a link to the PDF of the memo which says: "The purpose of this memorandum, issued pursuant to authorities under" - and I'm only going to read the introduction to it because it's a multipage memo. Actually, I have two pieces of it. But "...under section 3553(b) of Title 44, U.S. Code, and Title XXII of the Homeland Security Act of 2002, as amended, is to remind agencies of their legal requirements to use Einstein 3 Accelerated's Domain Name System" - that's Einstein 3A, as it's referred to henceforth - "the Domain Name System (DNS) sinkholing capability for DNS resolution, and to provide awareness about" - as the memo's intent - "to provide awareness about recent security and privacy enhancements to DNS resolution protocols, in particular DNS over HTTPS and DNS over TLS."

Okay. So this Einstein 3 thing started off life as an intrusion detection system designed by the DHS's US-CERT. Version 1 of Einstein allowed the agency to - and we're talking on a federal networking level - to monitor traffic across all government networks. Version 2 of Einstein added the ability to spot suspicious traffic. So they're doing intrusion monitoring. And in Version 3, where we are today, with Einstein 3 Accelerated, went still further, preventing unwanted intrusions by known bad actors. So, oh, it added filtering. It offers useful DHS-specific services like sinkholing that override the public DNS records by blocking access to destinations that the DHS knows to be malicious. So they're doing what a lot of us have already had, their commercial DNS services and so forth.

Anyway, it also lets the DHS examine all DNS requests made by government users. And we know how useful that can be. The memo notes that, they say: "CISA encourages efforts to make network communications encrypted by default. Doing so increases user security, making it harder for attackers to monitor and modify communications. DoH and DoT add desirable security features to DNS resolution; however, federal agencies that use DNS resolvers other than E3A lose the protection that defensive DNS filtering provides, and E3A does not currently offer encrypted DNS resolution. CISA intends to offer a DNS resolution service that supports DoH and DoT in time." But not in time for this memo. "Until then, agencies must use E3A for DNS resolution."

And then, finally: "Required Action. In accordance with 6 U.S.C. 663 note, 'Agency Responsibilities,' ensure local DNS recursive resolvers use E3A as their primary and ultimate upstream DNS resolvers." In other words, we love DNS encryption, but not if we cannot monitor it for your own safety, of course; and not if we are unable to apply our spiffy Einstein DNS filtering bad domain sinkhole system to keep all you government workers from going to naughty places by mistake. And, unfortunately, we don't currently offer DoH or DoT encryption of our value-added DNS, so be super sure to turn off that new on-by-default DoH resolution that Firefox and Google and soon Microsoft will all be using. Or you're going to be in trouble.

Leo: Well, you can understand why; right?

Steve: Yeah, completely.

Leo: They want to control it. It's the same reason the British government called it, what is it, Public Enemy Number One or something.

Steve: That's right.

Leo: Right?

Steve: Yes. They want to control it. They want to be able to look at it. They recognize its value. And of course the problem is, if somebody wasn't paying attention, Firefox will update, and Chrome will update.

Leo: Right, and it's turned on by default.

Steve: And it's on by default. And so suddenly their DNS server stops getting any requests. And it's like, hey, wait a minute. Where did everybody go?

Leo: Aw. Maybe that's why it took so long and still takes long for DNSSEC to penetrate. People actually don't want DNS security.

Steve: It's a mixed blessing.

Leo: Yeah.

Steve: So our somewhat nutty, but endlessly clever, guys at the Cyber Security Research Center of the Ben-Gurion University of the Negev in Israel, specifically this time a Dr. Mordechai Guri, have come up with yet another sneaky, if not entirely practical, means of exfiltrating an air-gapped computer's digital data by way of its switching power supply.

Leo: Wow.

Steve: Mordechai's research - yup.

Leo: Wow.

Steve: Mordechai's research has been published with the title "POWER-SUPPLaY: Leaking Data from Air-Gapped Systems by Turning the Power Supplies into Speakers." I'll just share the abstract because that gets enough of what they have done. So he writes in his abstract: "It is known that attackers can exfiltrate data from air-gapped computers through their speakers via sonic and ultrasonic waves. To eliminate the threat of such acoustic covert channels in sensitive systems, audio hardware can be disabled, and the use of loudspeakers can be strictly forbidden. Such audio-less systems are considered to be audio-gapped, and hence immune to acoustic covert channels.

"In this paper, we introduce a technique that enables attackers to leak data acoustically from air-gapped and audio-gapped systems. Our developed malware can exploit the computer power supply unit (PSU) to play sounds and use it" - I know.

Leo: This is great. I love it. It's amazing.

Steve: "...as an out-of-band secondary speaker with limited capabilities. The malicious code manipulates the internal switching frequency of the switching power supply and hence controls the sound waveforms generated from its capacitors and transformers."

Leo: I've been telling you for years my power supply's been talking to me.

Steve: That's right. It's making those little squeaky sounds.

Leo: Hysterical. That's hysterical.

Steve: "Our technique enables producing audio tones in a frequency band of 0-24kHz and playing audio streams, i.e., WAV files, from a computer power supply without the need for audio hardware or speakers. Binary data - files, key logging, encryption keys, et cetera - can be modulated over the acoustic signals and sent to a nearby receiver, for example, a smartphone."

Leo: I bet the bit rate's not super high.

Steve: Not high, but not bad. "We show that our technique works with various types of systems: PC workstations and servers, as well as embedded systems and IoT devices that have no audio hardware at all. We provide technical background and discuss implementation details such as signal generation and data modulation. We show that the POWER-SUPPLaY code can operate from an ordinary user-mode process and doesn't need any hardware access or special privileges. Our evaluation shows that POWER-SUPPLaY sensitive data can be exfiltrated from air-gapped and audio-gapped systems from a distance of five meters at a maximum bit rate of 50 bits per second."

Leo: Oh, well, that's pretty good.

Steve: That's not bad. "We should remember that while no one is going to transfer anything massive from a computer at 50 bits per second, elliptic curves, which are coming into increasing use, provide their state-of-the-art security using only a 256-bit key which could be sent in five seconds. Since every bit of a key is critical, and no bits can be inferred from others, I would encode the burst transmission with ample error correction. That might add another 30% to its size, but only a couple of additional seconds for burst duration."

So what these guys do is they chose four different frequencies, and so they send two bits at a time using one of four frequencies, since one of four gives you two bits. And that allows them enough time to set it up. Essentially, they are drawing varying amounts of power. And in a switching power supply, the power supply achieves its efficiency by changing the speed at which it switches power from the mains into its down-regulated DC, based on how much power is being drawn. Draw more power, the switching rates increases in order to couple more power from the main side down into the DC converted side. So, yes, you do get a shift in tone based on how much power is being drawn.

So these guys set up four different power levels, found four different discrete tones that would be generated, and then take two bits at a time from the data they want to send, use those two bits to choose one of four power draws. The switching power supply changes its switching frequency. A smartphone up to five meters away is able to detect that difference and essentially demodulate those two bits back to what they were before.

Leo: Thomas Edison is saying, "I told you. I told you. DC's the only way."

Steve: Yup. So we've had an authorization bypass in something known as SaltStack. We've never had the occasion to touch on SaltStack before, but a serious security vulnerability in this widely used cloud resource management system changes that. It lives over on GitHub, although there is a commercial enterprise that essentially is responsible for it and deploys it and maintains it. It describes itself as: "SaltStack makes software for complex systems management at scale. SaltStack is the company that created and maintains the Salt Open project and develops and sells SaltStack Enterprise software, services, and support. Easy enough to get running in minutes, scalable enough to manage tens of thousands of servers, and fast enough to communicate with them in seconds."

They said: "Salt is a new approach to infrastructure management built on a dynamic communication bus. Salt can be used for data-driven orchestration, remote execution for any infrastructure, configuration management for any app stack, and much more. Salt Open is tested and packaged to run on CentOS, Debian, Red Hat Enterprise, Ubuntu, and Windows."

So the good news is we've never had the occasion to talk about it before. The bad news is now we do. F-Secure wrote that they discovered a number of vulnerabilities in the Salt management framework. And the route of this is a little interesting. So they explained, they said: "The open source Salt project is at the heart of SaltStack the company's product offerings, but is also very popular as a configuration tool to manage servers in data centers and cloud environments."

They said: "Salt is used to monitor and update the state of servers. Each server runs an agent called" - and I love this - "a 'minion' which connects to a 'master,' a Salt installation that collects state reports from minions and publishes update messages that minions can act on," is the master. "Typically," they wrote, "such messages are updates to the configuration of a selection of servers, but they can also be used to run the same command in parallel over multiple, even all, managed systems asynchronously."

So it's a communications infrastructure and layer with agents that run in cloud servers that connect to this master, and there's an established protocol that allows files to be transferred, commands to be issued and so forth. So it's very powerful in terms of the manipulation and authority that it gives the master over a potentially massive infrastructure. And so you can imagine security of this thing would be paramount. F-Secure says the default communication protocol in Salt is known as ZeroMQ, Z-E-R-O-M-Q. The master exposes two ZeroMQ instances, one called the "request server" where the minions can connect to report their status or the output of commands that they're given, and one called the "publish server" where the master publishes messages that the minions can connect and subscribe to.

"The vulnerabilities described in this advisory..." - and this was only recently published, I think it was late last week, yeah, like Friday. This is F-Secure saying "...allow an attacker who can connect to the request server port to bypass all authentication and authorization controls and publish arbitrary control messages, read and write files anywhere on the master server file system, and steal the secret key used to authenticate to the master as root. The impact is full remote command execution as root on both the master and all the minions that connect to it."

They said: "The vulnerabilities, allocated CVE IDs" - and there's two of them - "are of two different classes, one being authentication bypass where functionality was unintentionally exposed to unauthenticated network clients, the other being directory traversal where untrusted input, for example parameters in network requests, was not being sanitized correctly, allowing unconstrained access to the entire filesystem of the master server."

And I'll just stop for a second to say, Leo, this directory traversal thing, I mean, it was very clever with the notion of a hierarchical directory tree. I can't imagine how we would organize without it. But maybe the problem is that the ../.. thing...

Leo: It's terrible. It's really terrible.

Steve: It's just been such a problem for us historically.

Leo: Is this the same thing? It's that dot dot again?

Steve: Yeah, it's the ../..; it just keeps biting us.

Leo: It doesn't need to be hierarchical. There has to be a better way to organize information. Hierarchical is silly. Isn't it?

Steve: I don't know. I'm a big outliner. I mean, I organize everything in outlines.

Leo: You could use tags. You could use other taxonomies. Think of it as a database as opposed to an outline. Right?

Steve: Yeah. And so you use context in order to access instead of - yeah.

Leo: I mean, I just think there's a better way. We just do it because we've always done it that way.

Steve: Yeah. And unfortunately ../.. allows you to go back up and then come down, descend down a different branch.

Leo: Every time I type that, and I do it a lot because I do a lot of command line stuff in Linux, I just go, there's got to be a better way. And most, by the way, most Linux shells have some shortcut method so you don't have to do that. In many shells you just type the folder, and it just goes there.

Steve: Yeah, and often, because I'm still old school, I'll go wait a minute, ../../..? Or do I want one more ../..?

Leo: Right, yes, yes, it's cuckoo.

Steve: Did I go back far enough yet?

Leo: It's cuckoo.

Steve: If not, I'll just kind of go, oh, that's probably enough. And if not, then I go, okay, I got close. And then I go back another layer or two.

Leo: No, that's nuts.

Steve: Yeah.

Leo: It's a form of weird skeuomorphism, is what it is, because we're so tied to this hierarchy of folders, folder within a folder. But they're not within anything. That's just so humans can kind of get it. I don't know, I don't think it has to be an outline. Maybe.

Steve: So in any event, we'll fix this after the podcast.

Leo: Yes, shall we? Oh, good.

Steve: Yeah. "SaltStack engineers patched" - this is F-Secure - "patched these vulnerabilities in release 3000.2." Hopefully these have not been sequentially numbered releases. "And users of Salt are encouraged to make sure that their installs are configured to automatically pull updates from SaltStack's repository server" - nice that there is such a thing, that's cool. And then anyway, so repo.saltstack.com. A patch release for the previous major release version is also available, with version number 2019.2.4.

So anyway, I'm going to skip the rest of this. Oh, except that to note - oh, yeah. Actually, I can't. They finish, saying: "Adding network security controls that restrict access to the Salt master which listens on ports 4505 and 4506, being the defaults, to known minions" - that is restrict access so that only the known minions have access, or at least block the wider Internet - "would also be prudent as the authentication and authorization controls provided by Salt are not currently robust enough to be exposed to hostile networks." In other words, this is on the Internet, and those ports were open, and there was, as soon as this became known, there was no authorization or authentication in place.

And they finish, saying: "A scan revealed over 6,000 instances of this service exposed to the public Internet." And remember, that's 6,000 masters that are probably managing relatively significant cloud infrastructures that all have lots of minions calling in to see, you know, for instructions. So, I mean, it's just like you couldn't make a better botnet-y thing if you tried to. And here it was just like, oh, yeah, whoops. We have an authentication bypass. Yeah.

Okay. So I'm going to skip over the timeline except to say that there was a little bit of comedy here because they first - F-Secure first tried to notify SaltStack on March 12th. They said in their timeline: "2020-03-12: The GPG key for the SaltStack security team published on SaltStack.com had expired in 2018, and a request for an updated key was sent." The point being F-Secure needed a valid GPG key in order to securely transmit this horrific finding that they had to get into SaltStack's hands to say, guys, you've got a really serious problem here. But the GPG key had expired in 2018.

Now, I guess it's a good thing that it had been since 2018, apparently, that anybody had tried to use the GPG key. So they're not getting lots of security reports. Or maybe people are not, you know, calling them up on the phone and saying, hey. But F-Secure tried to, you know, did it the right way. But oops.

So they requested a GPG key, waited four days until March 16th. Repeated their request. Finally got a re-signed key on the contact page so was able to send a full vulnerability report to the SaltStack security team. Four days go by. Nothing. So they request confirmation of the receipt to the SaltStack security team. Four days go by. The SaltStack security team finally confirms receipt of the vulnerability report and that they are reviewing it. Yeah, I hope so.

Next event is April 7th. SaltStack asks if F-Secure has requested CVEs for the reported vulnerabilities, and F-Secure replies in the negative, recommending that SaltStack proceed to contact Mitre in order to reserve some CVE IDs. Now, what, eight days go by. F-Secure informs SaltStack that an Internet-wide scan - oh, by the way, guys - turned up over 6,000 publicly exposed Salt masters and expresses concern that these will be at risk of compromise when the vulnerabilities are disclosed. Also F-Secure requests information on SaltStack's plan for distributing fixes.

Now I guess they've got SaltStack's attention because only in the next day SaltStack informs F-Secure that fixes are being tested and planned for release "early next week." That's in quotes. I guess that's exactly what they said. F-Secure reiterates concerns over the number of Salt masters exposed to the public Internet and requests information about how SaltStack plans to communicate this release to their customers. The point being, like, whispering would be good at this point.

Two days go by. SaltStack says to F-Secure, informs them of their communication plans and requests the list of identified IP addresses that expose a Salt master to the public Internet. I guess that's reasonable. I mean, you could get your own, or you could ask the people who already have one. What the heck. As well as suggestions for alternative approaches to disclosure. Oh, that's good, like we've never had something bad like this, this big, happen before. That's good. So how would you suggest we got about this? That was on April 20th.

On the 23rd, SaltStack publishes advance notice to their users urging them not to expose Salt masters to the Internet - yo, that'd be good - and to prepare to apply the patch once it is published on the 29th. So that was the 23rd. So they get six days' notice. On the other hand, the problem is presumably these SaltStack masters, the 6,000-plus of them, are exposed because the minions need access to them over the Internet. And they haven't taken the measure, the security measure of doing IP, like minion IP address filtering that would only allow the minions on the public Internet to see, to have access to the publicly exposed master. Anyway, six days go by.

On the 27th, F-Secure requests information from SaltStack about the CVE IDs allocated for the vulnerabilities. A few days go by. F-Secure reiterates the request for CVEs to SaltStack. Same day, SaltStack responds with the allocated CVE identifiers. Same day, SaltStack publishes v3000.2 and 2019.2.4 addressing these issues. So let's see, that's on April 29th. The "early next week" - oh, it's in quotes because that was said on April 16th, and it didn't actually happen until the 29th. So, what, 13 days later. The following day, F-Secure publishes their advisory.

Immediately upon publishing their security advisory, F-Secure was asked by the tech press, who was keeping track of these things, what F-Secure expected to see next. They said they expect to see attacks in the wild shortly. Quoting, they said: "We expect that any competent hacker will be able to create 100% reliable exploits for these issues in under 24 hours." They cited the "reliability and simplicity" of exploitation.

And dropping the other shoe, sure enough, hackers wasted no time exploiting vulnerable Salt instances used in various infrastructures for server management and automation. Among the organizations that announced an intrusion, and these are of course the ones being responsible and saying whoops: LineageOS, which is, I was unfamiliar with them before, but I guess an Android-based OS company.

Leo: Yeah.

Steve: Vates, the operators of open source Xen Orchestra; the Ghost blogging platform; and even my very favorite certificate authority, DigiCert. By now hundreds of servers - both masters and clients, the minions - if not thousands, have likely been compromised. Because exploit code is trivial to create, F-Secure published nothing in order to protect companies that would be slow to patch. Unfortunately, not that that's going to be much help. However, several versions and proof of concepts were immediately made public. So it's likely that any still unpatched servers are toast. I've got links to four proof of concepts that are all public on GitHub in the show notes. And this happened just over the weekend.

In the few days since the attacks began, more than 134 messages have been posted to Salt's bug page. And they make for some sobering reading. I've got the link to them in the show notes, for anyone who's interested. But, I mean, it's like our entire infrastructure is down and has been hacked. The good news is, for whatever reason, all that the attackers, at least the initial attackers, seemed to be interested in doing was installing bitcoin miners, which seems like, well, I mean, I'm glad that that's the case.

On Sunday Jeremy Rowley, who's DigiCert's Executive VP of Product, posted the news to the Certificate Transparency Group which is hosted at Google. He said: "Hey all. I'm sad to report that we discovered today that CT Log 2's key used to sign SCTs" - that's Secure Certificate Transparency logs - "was compromised last night at 7:00 p.m. via the Salt vulnerability." And he posts a link to a story about this on Threatpost.

And he said: "All other DigiCert Certificate Transparency logs are unaffected as they run on separate infrastructure. We are pulling the log into read-only mode now. Although we don't think the key was used to sign SCTs," he says, "(the attacker doesn't seem to realize that they gained access to the keys and were running other services on the infrastructure)," he says, "any SCTs provided from that log after 7:00 p.m. MST yesterday are suspect. The log should be pulled from the trusted log list." He says: "Happy to answer any questions about what happened, the infrastructure running the other logs, or what remediation we are taking."

So that's all you could possibly ask for. Immediately going public, dealing with the problem, telling people we don't have any evidence, we don't believe that the key that an attacker technically had access to was used. But obviously they rotated the key. They're pulling all the certificate transparency logs since the time that they know that the intrusion occurred. And they will then replace them signed under a new key, and we're back to where we were before. We don't know about what other entities may have been attacked.

But this was an interesting instance, another example, I guess, of how quickly bad guys will now jump on and exploit, even if not apparently the sharpest knife in the box bad guys, I mean, they could have done way more than they did. They just immediately installed cryptominers. Many of the bug reports that are in that log of intrusions just talk about cryptomining being installed. So it's like, well, again, we sort of got lucky with this one.

What made this worse was that it was so easy to exploit. For example, comparing to the RDP bug from a few months back, where it was initially a crash, and people were not sure whether you could actually leverage it into remote code execution, then it was. We saw something similar with the SMB failure a few months after that, where it can crash things. So if it really takes skill to leverage these, that buys a lot of time for the systems to get patched. If it's something that someone stumbles on, and it is just dead simple, even with SaltStack notifying their customers to get ready for this - and we've seen that, for example, in high-profile instances where the publisher will notify their users, we've got something coming that you really need to jump on because we're afraid the bad guys are going to see what this is and jump on it just as quickly as you do. So get ready to receive something.

And then F-Secure, being responsible, unlike some recent hackers who weren't, by withholding all details from their disclosure, just saying we found something, there's a couple CVEs, and we hope everybody patches this thing immediately. So bravo to them for being responsible.

I wanted to quickly note, mostly just to make sure people got the news, that Adobe had a big non-Patch Tuesday Patch Tuesday last week. Or non-Patch Tuesday Tuesday. They released emergency updates for three of their widely used products that patch dozens of newly discovered critical vulnerabilities. The affected software was Adobe's famous Adobe Illustrator, Bridge, and the Magento eCommerce platform, containing a total of 35 vulnerabilities where each one of them is affected with multiple critical arbitrary code execution flaws. So, yikes.

And it's unclear to me, as I was thinking about this, how Adobe Illustrator 2020, a quite capable drawing tool used by millions of artists around the world, could contain five critical remote code execution vulnerabilities. It's a drawing program. But these days no one is ever content to leave anything alone. It's a race to add features. So I guess I wouldn't be surprised if it, who knows, uses UPnP to open a port to itself so it can communicate with the cloud for some reason. And if you're not careful, Adobe Illustrator will get compromised. Or maybe it's just an image rendering problem; and if you made the mistake of opening a hostile file, then yeah, you could get taken over. In which case it wouldn't seem that critical to me. But still. For what it's worth, make sure you're current.

Mozilla announced another welcome service. Remember how we have Send from Mozilla, a painless, unlimited use, locally encrypted, TNO large file transfer facility that unregistered users can use for file transfers of up to a gig, and registered users can use for file transfers of up to 2.5 gig. It's available for free at send.firefox.com.

Okay. So Mozilla will be addressing another constant source of annoyance with a new service called Firefox Private Relay. It is a one-click email alias creation service. When Apple announced a similar forthcoming service as part of their sign-in with Apple at the 2019 WWDC, I remember thinking it was a cool idea, and I was a bit envious, since it wasn't clear how easy it would be for me to use it since I'm using email over on Windows. But Mozilla's forthcoming Firefox extension may be just the ticket.

The extension will generate unique aliases on the fly whenever you just need an address, but really don't want to give out your actual primary address. I know that we all maintain throwaway accounts for that. But this is better, since it offers a nice built-in email alias UI to manage them. So, for example, if you reused one, and I would tend to, but it started to get spammed, you could disable it, or you could delete it. So the forthcoming service has entered testing last month in beta, which is currently closed. A public beta is currently scheduled for sometime later this year.

And Mozilla explained: "We will forward emails from the alias to your real inbox. If any alias started to receive emails you don't want, you can disable it or delete it completely." So we don't have it yet, but I think it's neat that we'll have one over on the Firefox platform, which would make it actually universal platform neutral. And that'd be cool. I wish Apple's stuff was not so hostile to everything but theirs. I mean, I'd love to be able to send an iMessage from Windows, but I haven't figured out how to do that.

Leo: No, can't do it.

Steve: No. Okay, Leo. You're not going to believe this one. I mean, really. You're not. Political correctness hits cybersecurity.

Leo: Oh, no. That's bad. That's a bad start.

Steve: For the most part I feel young and wonderful. But I think that I must be getting old and crotchety, since things seem to be getting increasingly weird. Get this. The U.K. government's cybersecurity agency wrote last week that it would stop using the terms "whitelist" and "blacklist" due to stigma and racial stereotyping surrounding - I'm not kidding.

Leo: But that's reasonable because just because it's black doesn't mean it's bad or white because it's good. I understand what they're saying.

Steve: Yeah, well, I mean, I understand.

Leo: It's like master and slave on the SATA chain. It's like, yeah.

Steve: Oh, can't have that anymore, no. No, I mean, I understand it. I just think, really? They said instead the U.K. National Cyber Security Center said that going forward it would use the terms "allow list"...

Leo: Yeah, and block list.

Steve: ...and "deny list."

Leo: Or deny list. That's okay, yeah.

Steve: Instead of the other two.

Leo: That's even clearer than "whitelist" and "blacklist," to be honest.

Steve: That's true, it is clearer.

Leo: Allow and deny, yeah.

Steve: Emma W., who heads up Advice and Guidance at the NCSC, wrote: "It's fairly common to say whitelisting and blacklisting to describe desirable and undesirable things in cybersecurity. However, there's an issue with the terminology. It only makes sense if you equate white with 'good, permitted, and safe' and black with 'bad, dangerous, and forbidden.' There are some obvious problems with this," she writes. "So in the name of helping to stamp out racism in cybersecurity, we will avoid this casually pejorative wording on our website in the future." So I guess, what, are we going to have "allow list" hackers and "deny list" hackers?

Leo: Allow hat and deny hat.

Steve: Allow hat and deny hat.

Leo: You know, I'm not against this. I think this makes sense because it is institutional racism that's been incorporated for so long we're just used to it. But if you think about it, white and black.

Steve: But no one is thinking about African Americans when you say "blacklist."

Leo: Well, no, but that's where that comes from is black is bad and white is good. That's where it comes from.

Steve: Black is dark. It's the absence of light. It's, you know...

Leo: I think allow, yeah, I think allow/deny is better anyway. There are some issues. I mean, what are you going to use, white out? Are you going to use deny out? I don't know. Allow out? There are some places where it's going to be hard.

Steve: So you're saying we're not going to remove the two colors from the vocabulary.

Leo: No. But to use "black" as pejorative.

Steve: How about greenlist, redlist?

Leo: Yeah, you could do that. Like lights, traffic lights. And as long as you don't get any green people, we're okay.

Steve: Yeah, but we have American Indians, so...

Leo: They're red, so - no, they're not really red. But nobody's really white or black or red.

Steve: Exactly. Exactly. Except those two guys on Star Trek that were like cut down the middle; remember?

Leo: Right, half and half, yeah.

Steve: And that was actually a fabulous racism statement because Kirk and McCoy, they were standing there, like why are you upset with each other? You're the same. And they said - they, like, looked at them like they were crazy. And they said...

Leo: We're not the same.

Steve: ...what are you talking about? He's white on the right. I'm white on the left.

Leo: It's perfect. It's exactly right.

Steve: You're kidding me, yeah.

Leo: We could do good hat, bad hat. Again, clearer than white hat, black hat.

Steve: A bad hat.

Leo: Bad hat.

Steve: I know it would be. But, I mean, I can understand the sensibility, I guess.

Leo: I don't think that's so bad.

Steve: So anybody making bad Chrome extensions is being put on notice. And boy, are they being given time. I thought, why is the deadline August 27th? So I looked at where we are today and where we are then. It's like, what? How did they come up with August 27th? It's a Thursday, like in the end of August. Like, what? It's 16 weeks from now. I counted. And it's like, okay, well, nobody wants to accuse Google of not giving anyone any notice. So it's way past time to do this.

Last Wednesday Google announced new rules for the Chrome Web Store which should cut down the number, like a lot, the number of shady Chrome extensions submitted and listed. They explained that due to Chrome's success as today's top web browser platform - yes, we know, by a wide margin - the Chrome Web Store has seen an influx of spammers and fraudsters. Google says that these malicious entities have been behind a rising number of duplicate, spammy, and purely malicious extensions that are now poisoning and drowning the Chrome Web Store in low-quality content.

So be warned. Get ready. Mark it on your calendar because it's a ways away from now, starting on Thursday, August 27th. Google will begin enforcing a welcome new set of rules. I think they should do it tomorrow, but no. Which will result in a large number of extensions being immediately delisted. These rules are meant to crack down on a series of practices which extension "developers," and I put "developers" in quotes because it's just junk, have been recently employing to flood the Web Store with shady extensions or boost install counts for low-quality content.

So we've got six things that are no-nos, coming not soon to a Chrome Web Store near you. Developers cannot submit extensions, for example, wallpaper extensions, that have different names, but provide the user with the same wallpaper. So they're just redundant nonsense. Can't do that anymore, guys. Sorry. Extensions are not allowed to use keyword spam techniques.

Leo: Good.

Steve: Yes, to flood metadata fields with multiple terms and have that extension listed across multiple categories to improve the extension's visibility in search results. So no more search results spamming. Developers are not allowed to use misleading, improperly formatted, non-descriptive, irrelevant, excessive, or inappropriate metadata. Well, we're going to be the metadata police, thank god. "Extension metadata needs to be accurate, and Google intends to be strict about it," they wrote. I don't know how they're going to pull it off, but turn some AI loose on it, please.

Fourth, developers are now forbidden - well, not now, but after late August - forbidden from inflating product ratings, reviews, or install counts by illegitimate means such as fraudulent or paid downloads, reviews, and ratings. Extensions that have only one purpose, and a dumb purpose, such as launching a web page or an app, will no longer be allowed. Yeah, just launch it yourself. Extensions that abuse browser notifications to spam users with ads or other messages will also be banned. Good. Good. Do it now.

Anyway, on that Thursday, August 27th - and we'll let everybody know when it's only two days away because that'll be useful, and that'll be on a Tuesday - Google says it intends to take down every extension that violates these rules. Once upon a time it was bragging how many they had, and the more the better because it made Chrome look good. Now, not so much. Once that happens, many thousands of junk Chrome extensions will disappear, poof, from the Chrome Web Store, and nobody will miss them. This will make, they said, searching for useful content on the site easier and safer than it has recently become. And they finish, saying the Chrome Web Store currently lists more than 200,000 extensions - 200,000. And most of them are the same wallpaper.

Leo: I think that's true.

Steve: God. Look at this graph, Leo. Warning about RDP is not crying wolf. Just as the Shodan Internet-wide application search engine saw a 41% upward jump in RDP endpoints appearing at the beginning of last month, Kaspersky has had their eye on the Internet and has seen a somewhat startling increase in scanning, RDP scanning, and brute force, what we're now calling "credential stuffing" attacks. This chart is sobering, to say the least.

Leo: Are these origin countries? Or the places being attacked?

Steve: No, the target of these scans.

Leo: Targets. So the U.S. is big, of course, yeah.

Steve: Yes. Yeah. And of course Spain. Basically the large stay-at-home countries, interestingly enough.

Leo: Yeah. China's flat.

Steve: Yeah. From the beginning, well, China's probably where the scans are coming from.

Leo: Right.

Steve: Whereas from the beginning of the year we can see scans for open RDP ports, we're kind of purring along between maybe 150 to 200,000 per day, that all changed about the beginning of March. Many countries saw a quadrupling in attack rate, that is, within their borders. And at one point toward the end of March, Spain hit about 1.2 million attacks per day. Later, the U.S. even exceeded that by crossing the 1.4 million attacks per day threshold. Although overall there has been some ebb and flow, securing RDP from random Internet access is critical. I've often said there isn't a way to have it safely exposed. Please put it behind a VPN, if there's any way you can. And if not, really take advantage of LastPass and use a ridiculously long password that you can't remember, and you never want to type in, and have LastPass generate it and remember it for you. You just - you really don't want to be subject to brute force attack.

Okay. A fun bit of errata. Recall that last week we talked about RPKI and BGP. And I noted with apologies my utter inability to properly pronounce the name of one of the major players behind the work to secure BGP. His name is spelled J-O-B. That's his first name, which I think maybe Job?

Leo: I like Job. Yeah, I think it's what you said last time, Job, I like that.

Steve: But I'm not doing the last name.

Leo: Could just be Snijders.

Steve: You think so? S-N-I-J-D-E-R-S?

Leo: Yeah, could be.

Steve: Job Snijders? Anyway, I still have no idea how he pronounces it. But I do know that he was aware of my failed attempts.

Leo: Oh, dear.

Steve: He tweeted: "Hey, @SGgrc. Thanks for the shout-out. You did a great butcher job of my name, ha ha. Let me know if you want to talk more about RPKI and BGP in your show." So anyway, thank you very much.

Leo: But didn't send the right pronunciation.

Steve: No. That would have been fun. But anyway, if this becomes a thing, if it happens, we've got a contact. We may have him on the show and ask him himself how he pronounces his name. I very much appreciate the vital work he's doing.

Leo: Yes.

Steve: However he pronounces his name. Two notes, one thumbs up, one thumbs down. I don't remember if we talked about it during the show or before or after or when. But "Devs," Leo, it is a total of six hours and 49 minutes of run time, chopped up into eight episodes. So it's a little miniseries. It's produced by FX, and it's streaming on Hulu. Or you can do as I did, because I don't have Hulu. Although actually you can just get a Hulu subscription. You get a free trial for a week, and you could watch the whole thing for free on Hulu and then resign before you started to pay. I just bought it from Amazon for 13 bucks. Oh, my god. "Devs," D-E-V-S. And it is really good. I've not finished it yet. I've got about three episodes left.

Leo: See, one of the reasons I really liked it is it had a coherent ending to what must be very challenging. Well, you'll see when you get to the penultimate episode. You'll go, "I have no idea how they're going to solve this one."

Steve: Oh, neat. Anyway...

Leo: They paint themselves into a big corner, and then...

Steve: It is shockingly good and shockingly technically accurate. You teased me, noting about how in the opening conversation two of the main characters were sort of arguing about elliptic key versus RSA encryption.

Leo: I just loved hearing the words.

Steve: And the comparative merits of the two.

Leo: Yeah, it's awesome.

Steve: But, I mean, at one point a different hacker needs to securely look at a thumb drive, so he pops the keyboard off the top of his laptop and then, with a little screwdriver, pops two things off and then removes a little board and explains that he just removed the WiFi. And it's like, he used the little SMA connectors and popped the two antennas off of...

Leo: Isn't that cute.

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

Leo: Yeah, they've got a - whoever the technical advisor is, is actually knowledgeable. Yeah, it's good.

Steve: Yeah. I really like the characters. Everybody will recognize the lead guy's assistant. She played a big role in "Picard," the unfortunately not, I think, bound to be repeated CBS "Picard" Star Trek series. I ended up not being that impressed with it. Unfortunately, he was my favorite captain, but he's gotten pretty old. So he's not the Jean-Luc we know.

Leo: And you don't know who the long-haired guy with the beard is.

Steve: I don't, no.

Leo: But he's much better known than anybody else in the cast. It's Nick Offerman, who was Ron Swanson on "Parks and Recreation" and is beloved by the geek folks. He did a wonderful Yule Log last year in which he sat in front of a fire drinking whisky for eight hours. Nick Offerman's great. And this is a role very different from Ron Swanson or any other role you've seen him in. He's really good; isn't he?

Steve: I've never seen him before, and I am stunned by the quality of his performance. It was just, I mean, I was just like, I just can't believe how good it was. So I don't think I'm over-singing its praises. And I will balance that with Amazon Prime's "Tales from the Loop."

Leo: So disappointed that that wasn't good. I haven't watched it yet, but I got your text, "Don't watch it."

Steve: Yeah. It just - it was - it had potential. It had a couple interesting episodes. But it just never got off the ground. And some of the things were very obvious. Some of them were kind of like, okay. But anyway, just I can't highly recommend it.

Leo: And I'm curious, Amazon Prime has a new show by Greg Daniels, one of the creators of "The Office," called "Upload."

Steve: I've seen the previews.

Leo: Yeah, it's interesting. It's not great, but I'd be curious what you think of it. So when you get around to watching that, yeah.

Steve: I told everyone last week that I hoped to be able to report some first actual results from the first actual live testing of the hardware-level AHCI driver development for SpinRite, and I can. We now have drive enumeration and communication with every single system our testers have thrown at the code.

Leo: Nice.

Steve: And it's probably 50 machines, at least, maybe more. We're now working with Intel, AMD, ASMedia, Samsung, and Marvell AHCI chipsets, which is every chipset we've encountered so far. The newer Intel chipsets initially caused a glitch because they handled hardware interrupts a bit differently from everyone else. And the Marvell chips were just brought online yesterday morning before I began working to assemble this podcast. The trick with the Marvell chips was that they failed to implement two important status bits that are clearly part of the AHCI spec, but they don't have them. So I allowed support for those two status bits to be optional.

At this point, the new AHCI driver code is running on everything that all of our testers have thrown at it. So while I'm waiting for someone to dig up something that it doesn't work with, I'm going to reenable the support for all of the older interface standards - IDE, ATA, and AHCI when it's in legacy mode. We should then have enumeration and visibility for every drive on every system without any BIOS involvement. And because it's fun to see how fast our drives go, it's fun to see how long it will take SpinRite to run, and because people love benchmarks.

GRC's DNS Benchmark, I just looked because I was curious, is again our most often downloaded freeware. It's been downloaded a total of, yesterday when I looked, 5,620,883 times. And I don't double count for mass downloaders. That's a good, accurate count. And it's currently being downloaded at the rate of a little over 4,000 times per day, every single day. So not quite as many people as are showing up positive for COVID-19 in New York, but it's holding steady. Hopefully they'll get their count down.

Anyway, I want to verify our ability to transfer data. And since it'll be fun to know how long the next SpinRite will take to perform a problem-finding intolerant read scan of our new bigger drives, I'm going to - and actually I've already got the code written - perform a very accurate drive benchmark which will then allow everyone to test and compare all their drives' true performance. So this will end up being hopefully a popular and used and tested piece of freeware which all of the listeners on this podcast will have early access to because, again, if it doesn't work for someone, I want to know because that would mean that SpinRite would also not work. And that I really want to know. So anyway, making very good progress. And I will keep our listeners informed as we move forward. I might have it next week, who knows.

Leo: Wow.

Steve: Okay. So I guess the Internet is getting more explicitly political. China has proposed a wholesale revamp, which other authoritarian countries, including Iran, Russia, and Saudi Arabia, have vocally supported. So there's actually been dialogue about this. There's quite a lot of heated rhetoric surrounding the topic. And as with contact tracing, where it's complicated, the rhetoric and the technology might be quite different from one another.

So in working to get a feel for what was actually going on, I started with a PowerPoint slide presentation I found which had been, well, it had been given during one of the workshops and seminars conducted by the ITU, the International Telecommunications Union. And it was done by the main architect and mostly was sort of technological bullet points. It's not something anyone could implement anything from. But it's at exactly the right level to gain some overall understanding of sort of like what's the gist of this, which is my goal is to share that with our listeners.

So for their key network technology requirements they explained that they felt there was a need for more flexible, variable-length IP addresses to adapt to diverse scenarios and to be backward compatible with IPv4 and IPv6. They also support addressing and routing optimized based on communication entity semantics and something called "digital twin relationships." Some time was spent talking about "digital twins," but it was never very clear to me from the slides what they were. But they're also trying to support real-time what they described as "large-flux" communication in virtual physical fusion scenarios combining ubiquitous AI theory, whatever any of that means.

There's a lot of talk about the future need for truly massive bandwidth and guaranteed low latency in support of, and I'm not kidding, holographic transmission and applications such as a telesurgery. They also want secure, reliable, and resilient connection among massive heterogeneous networks - you know, why not? - and future network architecture supporting multi-ID space and these digital twin relationships.

So one thing that I think caused the hairs to stand up on the back of some people's necks was something that this system incorporates called "identity-based routing." That is, as opposed to IP-based routing. So it's one of the most interesting aspects of this, and I think it's, as I said, what I think frightens people. The slides explain: "Instead of mapping all information into network addresses, diverse IDs are used to indicate the destination which improves routing capabilities." In other words, this system proposes that it would be possible to address a connection to an individual based upon some sort of identity descriptor. And the network, necessarily knowing where this individual is, would route the connection to them.

And of course you can understand why that makes a lot of people uncomfortable, in order to do that, what you need to know, like where the person is. The slides explain: "Packets are addressed by semantic content metadata, of which old-style IP addresses become a subclass." And examples of semantic content are given: content ID, device ID, people ID, and service ID. So things have identifiers; and you can say I want this to go to this device or this person or this service, and somehow the network makes that happen. They suggest that this would enable network layer deterministic forwarding rather than our current "best effort" forwarding, which of course famously makes the Internet go today.

Leo: This is the antithesis of net neutrality.

Steve: Yes.

Leo: This is the exact opposite.

Steve: Yes.

Leo: All bits are not equal. All bits are inequal. And we can do what we want with them.

Steve: Yup. Some bits are better than others. And the overall global network would be redesigned with a sort of super QoS, that's my term, a super Quality of Service structure, which they assert will be needed to satisfy future scenarios. So the network's available transit bandwidth would be divided into slices, and they suggest some examples. An AR/VR slice would have low latency, less than 20 milliseconds. What they described as a self-driving slice, and I guess they mean autos, would have a latency much lower than that, even, of less than five milliseconds. And then a teleprotection slice would offer jitter of under 50 microseconds, so extremely low jitter. So the idea being you can somehow ask for different network characteristics and get that from the network.

And a lot of this they then pulled together on a slide titled "End-to-end communication requirements for intrinsic security," which is where they sort of are the most clear, and it's kind of chilling. They have authenticity where they say more than one third of autonomous domains right now, today, on today's Internet, do not have prevention mechanisms for IP address spoofing. And so they don't like that. Accountability versus privacy, they talk about the tradeoff. Exposing IP, port, and time to live on the wire decreases privacy, but anonymizing them decreases accountability.

So they say: "IP headers design should take into account a tradeoff between accountability and privacy." And you can kind of feel that they're willing to trade off accountability, well, they're willing to trade away privacy for increased accountability. For confidentiality and integrity, they say: "The current key exchange mechanism has many vulnerabilities." What they described as intrinsic identity keys would be used, not relying on third parties. But there's no notion of how they pull that off.

And then availability. "Avoid the unavailability of target network resources, computing resources, storage resources, et cetera, caused by DDoS attacks." They don't like those. "It is necessary to combine authenticity and accountability to build a multilevel verification filtering system for inter-domain and intra-domain traffic." So from this diagram, what is it, it's "Intrinsic Security for Privacy Protection in Future Networks." It's a diagram that's kind of like on a beige background behind it and a blue box in the lower left. Anyway, having looked at this diagram and studied it for a while - yup, that's the one, Leo.

So the system encrypts the sender's ID and the sender's local IP. They are then authenticated along with the hosting network. The traffic moves to the Autonomous System's edge router, and we know what that is from the last couple weeks when we were talking about BGP and autonomous systems like an ISP, and they have an Edge router that connects their whole subscriber big network to the Internet backbone. So the traffic moves to the autonomous system's edge router, which adds - and this is all new, of course, well, all of this is new, an ASID, an autonomous system identifier which is generated by an HMAC. The traffic traveling over the Internet contains this ASID verifier, an encrypted ID, and an encrypted IP.

As the packet attempts to enter the destination autonomous system, the sender's ASID, that is, the originating Autonomous System ID HMAC, is authenticated to permit it to enter. Then its content is made available, and the destination address pieces are decrypted. And what's most creepy is, in the flow chart diagram, there's this overriding auditing agent which is never described. It appears across the top with dotted red lines flowing down to all of the key elements on both sides. And in a box with a legend describing some of the abbreviations, the dotted red line is labeled "Shutoff Protocol," with not much else said, yeah. Apparently...

Leo: You need say no more, I think.

Steve: Apparently, if the auditing agent is unhappy, that traffic does not flow. There's also ample attention paid to the needs imposed by "holographic" communication which appears to literally mean the display of holograms. There's like a 3D body outline in one of the slides throughout the network. So they have ultra-high throughput, customizable priority and strategy, reduce complexity, I don't know how that happens, and...

Leo: Indeterminacy.

Steve: Thank you.

Leo: Indeterminate.

Steve: Indeterminacy. Reducing, we're reducing complexity and indeterminacy.

Leo: We don't want that.

Steve: We don't want lossy transmission to affect the quality of content. And basically they're changing everything. And inherent network awareness. I don't think they like our autonomous packet routing approach because it just - it's too indeterminate.

Leo: You have to realize that the Internet was designed exactly the opposite of this.

Steve: Yes.

Leo: For very specific and good reasons.

Steve: Yes.

Leo: And this just counters the whole thing.

Steve: Yes.

Leo: And who the hell is doing holographic communication anyway?

Steve: Yeah.

Leo: What are they talking about?

Steve: I really - there was a sense that they were trying to use these really inflated...

Leo: It's the future. It's the future.

Steve: ...future things. And so we're going to have to change the plumbing.

Leo: Yes.

Steve: In order to support our future holograms, our holographic dictator.

Leo: Thank god we have holographic communications now.

Steve: Yeah. So anyway, that's a sense for what they're aiming at. RIPE, which is the EU's Internet governance body, blogged about this proposal under the questioning title: "Do We Need a New IP?" That full blog post is long and winding, but the official response is short and to the point. Well, it's not quite as rude as saying no. But it's titled: "Response to 'New IP Shaping Future Network' Proposal." I've got the link to the PDF. It's only two pages.

It reads: "During this group's last meeting in September 2019 a number of proponents introduced a proposal titled 'New IP, Shaping Future Network,' proposing what was described as an opportunity for a strategic transformation of the Internet. The RIPE NCC appreciates the opportunity to respond to the proposal by means of this contribution. Since first being described in 1974, Internet protocol architecture has totally transformed our societies and economies, as the design philosophy of an open and flexible Internet that has allowed for an unprecedented number of life-changing innovations.

"It is true that, boosted by privatization and increased competition, these technological breakthroughs have also dramatically changed the telecommunications landscape. While it was once common to carry TCP/IP data streams across traditional PSTN infrastructure, those original telephone systems have now also evolved and are carried using the Internet protocol." In other words, it's sort of flipped it upside down. Instead of the TCP/IP being carried by the phone, now the phone is being carried by TCP/IP.

"Throughout its lifetime, the Internet protocol has also adapted to accommodate changing requirements and new technological insights. The most noticeable and impactful change was the redesign toward IPv6, with its 128-bit address space to overcome scalability issues and support growth beyond expectations. We believe it is exactly that open and adaptable nature, not only of the technical architecture but also the surrounding governance models, that is fundamental to the Internet's unprecedented success.

"Growing alongside its technical infrastructure, the Internet's governance model has also evolved to include new stakeholders and accommodate innovations, unforeseen use cases, and unimagined growth as the Internet became the fundamental technology powering our societies and economies. The open, inclusive, multi-stakeholder approach throughout its development, both in the technical forums that created its standards, as well as in the governance of its resources, has made the Internet what it is today.

"While we recognize the need for both the technical standards as well as these governance models to continue evolving, we strongly believe such evolution should take place from within the organizations and structures that invented the Internet and have supported its evolution throughout its history. We also strongly believe that any rationale for change must be carefully evaluated by all stakeholders in an open and transparent process to achieve consensus. The RIPE NCC is deeply concerned by what has been proposed here. We are especially concerned by the notion that this proposal represents an opportunity to steer away from the traditional bottom-up decision-making model.

"We also believe the technical rationale presented is flawed, and find the suggested alternative designs to be both unrealistic and unproven. Furthermore, if any of the proposed solutions could be developed to a mature and production-ready standard, market adoption is very uncertain and would take decades to accomplish. The RIPE NCC is of the opinion that the proposal is premature, and that following through with any of the suggested work would create significant overlap with the ongoing work of other stakeholders, in particular that of the Internet Engineering Task Force.

"Although some of the issues mentioned may warrant further study, we insist that any work on the evolution of the IP protocol layers and the associated technical standards be left to the Internet Engineering Task Force and be conducted under its governance. The RIPE NCC recommends that TSAG, and the ITU Telecommunication Standardization Sector in general, not make any of the suggested changes to its structure or pursue any work items related to this proposal that would evolve the Internet protocol stack under the ITU's remit." In other words, hands off.

So that's what they had to say, not surprisingly. This is interesting, I guess, from a sort of theoretical "what if" model; and it sure did ruffle, as I said, a bunch of feathers. But it's clear that those of us who have some deep affection for the operation of today's Internet, warts and all, have little to fear. Let's remember that we can't even get people to upgrade their version of TLS.

Leo: This would be a lot hairier. This would be a lot more, yeah.

Steve: Oh. Oh, Leo. It breaks everything.

Leo: Yeah.

Steve: It would scrap the entire investment in all of everyone's existing Internet networking equipment. And I was thinking about, okay, first of all, no one wants it. But if such a thing, I mean, maybe there would be a completely next-generation thing once the holograms are needing more bandwidth, demanding their own bandwidth. If such a thing should ever happen, first of all, it's a long ways off. And the only way I can see anything like that ever happening would be as an overlay of an entirely new next-generation network, which would initially run in parallel, and only in a few locations, alongside today's Internet. And it would gradually, maybe, over time, eventually replace the older Internet.

Which, come to think of it, is exactly the way the original Internet was born. It only connected a few things, and it didn't try to replace the telephone. And then it got bigger, and it connected more and more places, and then it got more and more use, and it just sort of happened. So I can imagine some next net that would happen sort of the same way. It would just, you know, some people would start connecting things up, and other people would go, hey, we need to get one of those boxes if we want to play with those holograms.

Leo: Consider the holograms, my friend. Consider the poor holograms. Steve Gibson, he's at GRC.com. That's his website. Now is the time to sign up for SpinRite because then you'll get the beta. Maybe even as soon as next week you'll get to try out all these new little bits that Steve's incorporating into the next generation of SpinRite. But only current SpinRite holders get to play with that stuff. GRC.com.

While you're there, get a copy of this show. Steve has 16Kb audio for bandwidth impaired. He also has 64Kb audio, even transcripts, which probably - those are just text, so I bet those are the smallest versions of the show of all. They're all at GRC.com. And there's lots of free stuff there while you're visiting, including ShieldsUP! and the DNS Benchmark, all of that. Steve's on Twitter at @SGgrc, and you can leave messages there for him, or on the website at GRC.com/feedback. DM him on Twitter. And that's a good way to stay in touch and keep up with what Steve's doing.

We do this show every Tuesday, right after MacBreak Weekly, so it's about 1:30 Pacific, 4:30 Eastern, 20:30 UTC. You can watch the live stream of our production at TWiT.tv/live. That goes day and night, audio and video streams there, TWiT.tv/live. You can get the show from the website, TWiT.tv/sn. It's on YouTube. You can ask your favorite voice assistant. "Hey, favorite voice assistant, play Security Now! podcast." It'll play the latest version.

I think that's all I have to tell you. Oh, subscribe. That's the one other thing I want to say. If you haven't subscribed, that's the best way to get it because then you'll get every episode the minute it's available, and you can listen to it at your leisure.

Steve, thank you. Have a great week. Stay healthy. Stay safe.

Steve: My friend, next week we will continue our 12-and-a-half-year trek through the security wilderness.

Leo: Holy moly. Thank you, Steve. We'll see you next time.

Steve: Okay, buddy. Bye.