I'm Not a Robot! (Really)

Description: This week, Leo and I discuss law enforcement and the Internet of Tattling things, a very worrisome new and widespread PHP eMail vulnerability, Paul and Mary Jo score a big concession from Microsoft, a six-year-old "hacker" makes the news, Apple discovers how difficult it is to make developers change, hyperventilation over Russian malware found on a power utility's laptop, the required length of high-entropy passwords, more pain for Netgear, an update on the just finalized v1.3 of TLS, the EFF's growing "Secure" messaging scorecard, a bunch of fun miscellany - and how does that "I'm not a robot" non-CAPTCHA checkbox CAPTCHA work?

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-593.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-593-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. First show of the new year, and of course there's lots of stuff to talk about, including the hacking of a Vermont power station and that Russian spyware that the Defense Department found. We'll talk about that. There's a whole lot more coming up. Security Now! is next. It's great to be here for the new year.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 593, recorded Tuesday, January 3rd, 2017: I'm Not a Robot, Really.

It's time for Security Now!, the show where we talk about security, now, right now, with Steve Gibson, the man in charge at GRC.com. He's the security guru. We talk a lot about privacy, too, and how things work. Hey, Steve. Happy New Year.

Steve Gibson: Leo, great to be with you for 2017.

Leo: Wow. Wow. Wow.

Steve: Yes. Yes, yes, yes. And I have a feeling it's going to be an interesting year of - actually, the first story is the Internet of Tattling things, so I...

Leo: Oh, I know what you're going to talk about. This is the criminal suit, right, or

whatever, the prosecution, yeah.

Steve: Yes, yeah, from Arkansas. So we have law enforcement and the Internet of Tattling things. A very worrisome new and widespread PHP email vulnerability which is present in nine million different websites at the moment. Paul and Mary Jo scored a big concession from Microsoft just before the end of the year on Windows Weekly. We have a six-year-old hacker who makes the news over the holidays. Apple also discovers how difficult it is to make developers change. Some hyperventilation over Russian malware found on a power utility's laptop and how the press is just going bonkers for, like, no apparent reason.

I found something really nice about the required length of high-entropy passwords from someone who really knows what he's talking about. More pain for Netgear. An update on the just-finalized v1.3 of TLS. Oh, and I realized I changed something here, but I didn't change it in the notes. The EFF, I found a new page on the EFF site which is an amazing sort of self-serve security reference that I know that our listeners are going to be interested in. We have a bunch of fun miscellany. And we're going to do what we couldn't get to last time, two weeks ago, because we just ran out of time. And that is, I'd also named this week's podcast: I'm Not a Robot, Really. How does that checkbox non-CAPTCHA CAPTCHA manage to be so simple and, apparently, effective? So I think another great podcast to kick off the new year.

Leo: I was curious if you would talk a little bit about - and I guess there's really not much to say - what the intelligence agencies released as their explanation of the Russian hacking. Dan Goodin had a great takedown on Ars Technica, and I'm sure you read that.

Steve: And actually it's our topic for next week.

Leo: Good. Good. Because there was a - I looked at it. I don't remember the name of the security researchers that went through it. And it looked as if it was nonsense, basically.

Steve: Well, yeah. It looks like there's, from what I saw, they're describing it as some old and obsolete Ukrainian malware.

Leo: PHP malware, which is widely available through the web.

Steve: Exactly. But the details of it I thought were interesting, so I thought...

Leo: Yes, yes.

Steve: And there's been a huge amount of interest from our listeners. So I said, okay, let's, you know, I will share everything that we know from a technical standpoint. And that's our topic next week.

Leo: Yeah, we can talk about it then. Because, you know, there's a meta story about it which is, if this is their evidence, it sucks.

Steve: Yes.

Leo: But there's also the probability that they can't really release their evidence because to do so would impair their operational capabilities. Maybe. I don't know.

Steve: Yeah. Well, and of course, as our listeners know, that's what stops me completely because...

Leo: We don't know.

Steve: Anybody can gossip, and you can get that anywhere.

Leo: Right.

Steve: But I'm not interested. We need details and facts.

Leo: Right. We don't know.

Steve: And it's inherently, they're inherently not going to be available. But at least in this one case we know everything about this particular piece of malware which itself looked pretty interesting.

Leo: Okay, good. Next week.

Steve: So we will do a deep dive next week.

Leo: All right.

Steve: You have 28 people; right?

Leo: Yeah.

Steve: And I had 23 at GRC's height. Well, or depth, depending upon how you do it.

Leo: Oh, god.

Steve: And one of the things that I noted, which you just alluded to, is it's remarkable how people are constantly having birthdays.

Leo: A lot of partying.

Steve: When you get a certain critical mass of people, it's someone's birthday constantly.

Leo: Almost all the time, yeah.

Steve: And it's like, wait a minute. Already we have another birthday? So obviously, if it was evenly distributed, and you had 24 people on a 12-month calendar, you'd have an average of two per month.

Leo: A month. That's about right.

Steve: So it's like, okay, that's every other week. It's like, okay, wait a minute. Again, another birthday?

Leo: That's a lot of cake.

Steve: Stop growing.

Leo: Now, you could, by the way, use ZipRecruiter to narrow down your candidates by birth date, and have them all have the same birthday, and really simplify things. But I wouldn't recommend it. If that's your top priority.

Steve: Now, I wonder if that would be considered hiring discrimination?

Leo: Yeah, I think it would, yeah.

Steve: Because, well, it's not something that they have any control over. So in general, things that are...

Leo: It has to do with age, so, yeah.

Steve: Yeah, yeah, yeah. Okay. So I've been thinking about your comment, about this

Picture of the Week really is a spoiler for Westworld. And you're right. So...

Leo: Close your eyes.

Steve: And here's the problem.

Leo: Yes.

Steve: The thing that makes this such a brilliant picture is why it's such a horrific spoiler.

Leo: Yeah, right. That's the humor of it.

Steve: Yeah, it's so good because it's not something you know unless you finish the series. So do not, listeners, look at the first page of the show notes. Do not even download the show notes. Forget about everything.

Leo: And if you [crosstalk] audio, you're safe.

Steve: Yes. And I don't know why you haven't finished Westworld. What's wrong with you?

Leo: Right, that's the real question.

Steve: Maybe you don't have an HBO subscription. You're waiting for it to appear on the torrent of your choice.

Leo: Right. I think that's true. No, I think a lot of people are going to - this happens with all the HBO shows. They sell them later. So a lot of people say, well, I don't have HBO, but I'll just buy it. So we don't want to spoil it for you. So if you've not seen Westworld, close your eyes. Actually, if you've seen just a few episodes, that's really when you need to close your eyes.

Steve: Yeah.

Leo: This is not revealed till later in the show.

Steve: Yeah.

Leo: But if you have seen it, I think you're going to find this very funny.

Steve: It's the best thing ever.

Leo: Oh, man. Is this mean. It's a CAPTCHA. It's a CAPTCHA with that checkbox that says "I am not a robot," and a picture, and a caption. That's the part that would be revealing if we said.

Steve: Yeah.

Leo: Somebody is not quite sure if he or she is a robot.

Steve: Exactly. And from what we know of Westworld, there does seem to be some confusion among the ranks.

Leo: Yes.

Steve: And that was one of the really interesting unknowns for a long time was what the robots believed about themselves. They had a back story which it was explained was put in place to help form, to like anchor their personality. And so the way they handled clear problems with reality was interesting because of course they weren't in reality, they were in an amusement park. So anyway.

So a man was found dead, unfortunately, in the hot tub, the backyard patio hot tub of the person who then became the murder suspect. And his body had cuts and bruises consistent with a fight, and there was blood in the hot tub water. The suspect has been charged with and will be tried on charges of first-degree murder. And this happened in Arkansas. Detectives in Arkansas then wished to obtain data from the suspect's various IoT systems, things that might provide some additional information. And this is really the anchor of the story that we'll get to in a second, but I want to sort of cover the facts.

So the thing that made the news wasn't what I think is a little more interesting to our audience, although the specifics are, and that is that this person, the suspect in his home had an Amazon Echo. And Amazon received a warrant, essentially, requesting any information, any audio that the Echo might have picked up and that Amazon had on their servers. And we'll talk about what Amazon said in a second. But essentially they've asked Amazon for anything that the suspect's Amazon Echo may have overheard. So...

Leo: Wow.

Steve: Yes.

Leo: We've talked about this before. I mean, it wouldn't have overheard anything unless by accident, as the guy is being murdered, he shouts, "Alexa." Right?

Steve: Right, yeah.

Leo: Okay.

Steve: You know, which edge of the knife do I use? And I'm not sure what you want to ask Alexa at a moment like that, but yes.

Leo: Now, I don't know if it's the same case, but there was also a subpoena for recordings of some other device.

Steve: Well, what they noted was interesting, was that this person, apparently he was into IoT, for better or for worse. He had a smart water meter. And it logged 140 gallons of water used between 1:00 a.m. and 3:00 a.m.

Leo: Aha.

Steve: The night the victim was found dead in the suspect's hot tub.

Leo: You see we have caught you. You are guilty.

Steve: And so the investigators assume, and they're alleging, that the water was used to wash away evidence of whatever it was that transpired on the patio.

Leo: Right, right.

Steve: Or maybe the hot tub was dumped and then refilled because it was too bloody? I mean, who knows? So the examination of the water meter and the request for stored Echo information raises a bigger question about privacy. At a time when we have any number of devices now tracking and automating our habits at home, the question arises: Should that information be available for use against us in criminal cases? Now, the defense attorney, of course, argues that it should not be, saying one has an expectation of privacy in one's own home. And, she says, "I have a big problem that law enforcement can use the technology that advances our quality of life against us." Well, okay. Oh, boohoo. I mean, her job is to take that position. But it does really raise the question.

And from our audience's standpoint, and people who follow the podcast, there's also, as we would well know, a real question of the reliability of information from smart home devices, since accuracy can, as we know, be an issue for any number of readily hackable IoT gadgets. So someone could plausibly be framed by their IoT devices because, you know, they're being hacked all the time. So it might not be your light bulb DDoSing some random dotcom site somewhere else. It could be your IoT device, your light bulb, basically tattling on you, saying, yes, he was home at 5:00 a.m. because I got turned on, when in fact that information could be planted externally.

Leo: But all of that would be - that's part of the court's, you know, and the

testimony, and you have to explain all that stuff.

Steve: Well, and we know, I mean, how many stories have been built around the whole chain of evidence issue, you know, the legal formalization of verifying that at no point was the chain broken, was something never subject to tampering, out of control of law enforcement and so forth.

Leo: That's right, yeah.

Steve: So an Amazon spokesperson - first of all, Amazon denied them access to any recorded audio.

Leo: That's interesting, too; isn't it.

Steve: Yes.

Leo: By the way, if you had access to the suspect's phone, and he had the Amazon app, the Echo app on his phone, all of the recordings are available through that app.

Steve: Yes, and we're about to get to that because Amazon does store recordings of what it picks up in the cloud. So the Amazon spokesperson denied the audio, but did provide purchasing information, that is, non-Alexa - oh, sorry, non-Echo.

Leo: That's all right, I already screwed it up by shouting. The jokes, by the way, the jokes here are endless on different commands that you could use that would be incriminating.

Steve: Oh, coming over the chatroom? Oh.

Leo: Yeah, the chatroom's been having a good time...

Steve: Quite prolific.

Leo: ...with this, yeah, yeah. I'll - go ahead. I won't interrupt.

Steve: Okay. So they did provide some information. Probably the warrant asked for everything. And they said, well, we'll give you some of what you've asked for, like when did he purchase this new set of knives, but we're not going to give you the audio from the Echo. So the Amazon spokesperson formally said: "Amazon will not release customer information without a valid and binding legal demand properly served on us. Amazon objects to overbroad or otherwise inappropriate demands as a matter of course."

Okay. So then some facts were made more clear as a consequence of research that followed this. First, as we have said, the Amazon Echo only captures audio and streams it to the cloud when the device hears the wake word "Alexa." And as we know, it may actually - it's probably buffering it, looking for that word in the buffer. So it probably, in the same way that Apple's little animated picture widget captures before and after you actually press the snapshot button, I would imagine that Alexa is picking up some context in case it's important beforehand.

So there is a gated capture and stream to the cloud which surrounds the engagement of the Echo. And as we know, a blue ring on the top of the device provides a visual indication that audio is being recorded. So the blue ring is meant to be connected to Alexa is listening actively and streaming to the cloud. Those clips, or "utterances," as the company calls them, are stored in the cloud until a customer deletes them either individually or all at once. When that's done, the utterances are permanently deleted. And just in case someone didn't know, the microphones on an Echo device can be manually turned off at any time.

So I think this represents Amazon having thought this through. I mean, they couldn't possibly handle the bandwidth of everybody's Echoes and Dots in the world sending, you know, spying on all of their homes. And nobody would want that happening. So I think they made a good tradeoff. They've made it clear what they're doing. They're no doubt capturing sound samples, both to learn about Echo's master and probably over time to perfect their audio. And also, for example, if you ask Alexa something that the system, the Echo - oh, okay, I give up.

Leo: By now everybody's muted their system.

Steve: If you ask her something that she doesn't understand, I'm sure that sends a ping back to the mothership, saying this snippet was not understood. And probably there's minions somewhere listening to those, going, ooh, what a great idea for, like, what she should be able to understand. So there's no doubt some serious metrics and analysis going on on the back end to make it even better. So that's a tradeoff. And I think Amazon has done a good job. But the broader question, of course, is not this specific, but more like the power meter, which is a perfect example of information which once upon a time would have been lost, which law enforcement would want to have, if they could. And now to a much greater degree there is, and we talk about it all the time, the privacy versus tradeoff agreement, essentially. I mean, the nature of that is, if you've got - what's the thermostat on the wall, the...

Leo: Nest.

Steve: The Nest.

Leo: There are a lot of them now.

Steve: Yes. Exactly. If you've got a smart thermostat, if you've got something like a home security system which is doing stuff, I mean, essentially many of these things, not all of them, but many of them are monitoring in nature. So if your lights are being

monitored, your heating and air conditioning, the presence of your motion in rooms and so on, and that information is obtainable retrospectively, then that's just something to remember. And there may be a very strong defense, which is, "Your Honor, everybody knows these systems are being hacked all the time. My client was framed. He wasn't home, even though his home says he was, because somebody arranged, you know, planted that information. And the prosecution is being fooled by believing that this could be trusted." So we live in very interesting times.

Leo: You know, it's going to be everywhere; right? This is all being gathered at all times.

Steve: Yup.

Leo: And law enforcement's going to want it. And you know what, if it convicts some murderers, I'm not sure I'm against that. I just worry that it'll be used against dissidents and others for political reasons as opposed to real crime.

Steve: Well, in Southern California we went through a phase, maybe about five years ago, where all intersections got cameras pointing down all four streets. And it's like, okay. And in fact it's funny because I remember looking when my particular intersection, because I'm on not a very busy street, got a signal. It used to just be stop signs, and people were getting killed, unfortunately. I think it took three, and then they said, okay, that qualifies for a stoplight.

Looking inside this thing, I mean, it looks like a chunk of the Level 3 server farm. It's an unbelievable amount of technology. And I'm looking at these lights that go reluctantly green, and red and yellow mostly, and thinking, it takes all this in order to run a stoplight? As an engineer I'm thinking, there's something wrong with this picture. But what apparently they have is this countywide incredibly high-bandwidth network that planned to have things like cameras pointing down every street in order to capture this information. So I don't know if it's being stored locally, or if it's being streamed, or what's going on. But, yeah, as you say, Leo, there are essentially monitoring posts being set up all over the place in order to provide services, but also to keep the peace. And as you say, let's hope it isn't abused.

So speaking of abuse, we have an extremely worrisome - PHPMailer is the very popular PHP library - zero day, meaning there is no fix for this, and this is very bad right now. It's a critical flaw which about nine million PHP-driven websites and many popular open source web applications are using worldwide, including WordPress, Drupal, 1CRM, SugarCRM, Yii, and Joomla, all use this PHPMailer library as their mailing agent interface. A Polish security researcher, Dawid Golunski, with Legal Hackers is the name of his group, discovered a critical vulnerability some time ago. And that vulnerability allows an attacker to remotely execute arbitrary code in the context of the web server and compromise the target web applications.

So this is a classic example that answers the question, how do bad guys get in to get up to their evil deeds in a website? It's exactly this kind of thing. I have the technical details, that we'll get to in a second, to sort of flesh that out. But in his posting he wrote: "To exploit the vulnerability, an attacker could target common website components such as contact or feedback forms, registration forms, password email resets and others that send out email with the help of a vulnerable version of the PHPMailer class." Okay, and

so that's, like, everything today is vulnerable. There is no fix for this. "A successful exploitation could let remote hackers gain access to the target server in the context of the web server account, which could lead to a full compromise of the web application."

He responsibly reported the vulnerability to the PHPMailer developers, who then patched the vulnerability in their new release, which was 5.2.18. And all versions before that release were affected. Turns out, however, they didn't fix it. And as a consequence of that patch, attention got focused on this, and the nature of the problem was published in a security mailing list. So the secret disclosure, essentially, as a consequence of the PHPMailer guys fumbling the fix, their fix was found not to have fixed the problem. The problem still persists, and now it's public knowledge. So Golunski immediately posted full details because the cat was out of the bag? Is that...

Leo: Cat's out of the bag.

Steve: Cat's out of the bag.

Leo: Why the cat was in the bag, no one knows.

Steve: Horses have left the barn...

Leo: Horses, yes.

Steve: ...[crosstalk] jumped out of the hat, whatever. Anyway, too late. So we have a situation where websites are vulnerable to deliberate exploitation. Okay. So the 2016 original patch was CVE-10033, which added, in version 5.2.17, sanitization of the $Sender variable by applying an escapeshellarg function to essentially protect the arguments before the value is passed to the mail function. Now, this is, unfortunately, this is a perfect example of sort of the equivalent of setting up a firewall to block only the ports that you know have ever been used for evil, otherwise leaving everything else open. So it's a blacklisting approach, rather than a whitelisting approach. Meaning in this context that there's a fundamental problem between the web interaction characters and HTML and the characters used by operating system shells.

And we run across this problem all the time. In fact, I had it. That problem I had that we talked about, what, six months ago? There was no way to exploit it for any mispurpose, but it was the one researcher that found a flaw in one of my forms where I hadn't - essentially exactly this. I was not perfectly removing special characters from showing them to the web page, which allowed a bad guy to carefully design something that would execute their own code in the page that the user received. You don't ever want to let that happen. As it turns out, there was not a way, due to all kinds of other things that I do differently than the rest of the world, to exploit that. And we talked about it at the time.

So this is one of those. This is where they just - they missed something for a long time, that there was a way to submit a form to get something that should have been protected from use onto the user's page. So in technical details it says this unescaped string gets passed to the mail function. However, what they fixed for that vulnerability does not take into account the clashing of the escapeshellarg function with internal escaping that is performed by the escapeshellcmd, which is performed by mail. So there it sounds like it's

- and I didn't dig into this infinitely because there wasn't any real purpose. It's going to get fixed quickly.

But on the fifth parameter there is interaction between those two. So as a result it's possible to inject an extra quote that does not get properly escaped, and break out of the escapeshellarg protection applied by the patch which was made in the first fix at 5.2.17. So an attacker could pass, and in this case it's the -X parameter of sendmail, to write out a log file with arbitrary PHP code and cause that to be executed. So what this means is that the current latest versions, both 5.2.19 and 5.2.18, are vulnerable to remote code execution despite the patch. And Dawid immediately posted all of that information and a proof-of-concept exploit because, why not, we now know how to do it. I mean, the truth had already escaped because they didn't fix it right the first time.

So if anybody feels they may be vulnerable to this, or know somebody who might, make sure you check to see if PHPMailer has been updated. It's going to take .20. That'll be the one which fixes it, 5.2.20. Anything previous to that has this problem. And it looks at the moment like it is very widespread.

So just before the holidays, Paul and Mary Jo had the Microsoft Chief Marketing Officer, Chris Capossela...

Leo: Capossela, yeah. He was on the show, yeah.

Steve: Right. And they got huge news coverage for Windows Weekly because in this interview, where they were going back and forth...

Leo: Something must have happened that I missed, because I was listening, and I was right there.

Steve: So, yeah, I mean, ExtremeTech, Techdirt, Business Insider, Softpedia, and many more all picked up on the story and referred to Paul and Mary Jo and Windows Weekly. What he said was, he said, "We know we want people to be running Windows 10 from a security perspective. But finding the right balance, where you're not stepping over the line of being too aggressive, is something we tried. And for a lot of the year I think we got it right." Okay. Many people would disagree.

He says: "But there was one particular moment in particular," he says, "you know, the red X in the dialog box..."

Leo: Yeah. He admitted that that was a mistake.

Steve: Right, "which typically means you can't" - means you, okay, I'm reading what he said, "typically means you cancel, didn't mean cancel." And he said: "And within a couple of hours of that hitting the world, with the listening systems we have, we knew that we had gone too far. And then, of course, it takes some time to roll out the update that changes that behavior."

Leo: Right.

Steve: "And those two weeks were pretty painful..."

Leo: Yeah.

Steve: "...and clearly a lowlight for us. We learned a lot from it, obviously."

Leo: And this is, by the way, why we love having Chris Capossela on because, unlike a lot of executives, he says the truth.

Steve: Yes, yes.

Leo: And you know what, that was valuable to me to hear Microsoft say, yeah, we blew it. People make mistakes. And what we don't really know is how long it takes them to fix it. Takes a couple of weeks, due to the nature of their systems or whatever.

Steve: Right. And the only pushback I would have about that is that the red X was like the last straw. That was after...

Leo: Yeah. No, no. Yeah.

Steve: ...six months of, I mean, even Paul was like, okay, this is, you know, like before the red X, remember when there was no obvious way, like either button you pushed caused the update? It was like, update now, or update later. Uh, wait a minute.

Leo: Right. That was dopey, too, I agree.

Steve: So, yeah. So their listening systems, I would argue, I mean, if they had some - what happened was they were willing to tolerate a very high level of annoyance, which finally did overcome even their high threshold, which they clearly were intending to tolerate in order to get Windows 10 deployed as far and wide as possible. And on this occasion, I just checked the Never10 page this morning, seeing 2.262 million downloads. And we've slowed down now to only 2,236 a day. So that gave people the option.

Leo: The other thing I would point out is we knew how loud the pain was because within our circle we heard it. But I think probably Microsoft's listening systems extend to the entire - remember, there's a billion and a half users of Windows - extend to that larger group, where I don't - I think the upset was probably considerably more muted among normal people; right?

Steve: And it's only the instances where the upgrade crashes or kills someone's machine that makes a big cloud, and everybody says, "Oh, my goodness." Whereas you don't hear from all the people who are like, oh, wow, this is better? Okay. And then they just go on.

Leo: We don't like it. But we're much more, I think - we're not typical. We're much more aware of how it was being thrust down our throats. And I think you also have to acknowledge that they did have - there was some merit in their idea of we'd like to get everybody on the same platform and then offer required but consistent free updates from that point on. That would really help, not only their business, but the security of the Internet in general. And so I understand how they were willing to push a little bit to do that. And I love it that Capossela acknowledged, hey, we pushed too far.

Steve: I agree. I agree.

Leo: And we realized that immediately.

Steve: Yes. It's nice to have someone from whom you don't get just total spin.

Leo: He's great. He was on last year at this time. And I think he'll be kind of a yearly thing. And I just - I want to encourage him because it is, it's unusual for a company to be forthright at all about this kind of stuff, you know.

Steve: Yup. Okay. So in this holiday season's classic demonstration of the security/convenience tradeoff, a six-year-old daughter uses her sleeping mother's thumb to purchase $250 worth of Pokemon toys for herself. "Ashlynd Howell of Little Rock, Arkansas is a precocious six year old. While her mom, Bethany, was sleeping on the couch, Ashlynd gently used her mom's thumb to unlock the Amazon app on her mother's phone. Ashlynd then proceeded to purchase $250 worth of Pokemon presents for herself. When her parents got 13 confirmation notices about the purchases, they of course thought that they'd either been hacked - well, technically they were, but not by someone remote - or that their daughter had ordered them by mistake."

Leo: Uh, not a mistake.

Steve: "But she proudly explained, 'No, Mommy. I was shopping.'"

Leo: I love it.

Steve: The Howells were able to return only four of the purchased items. So anyway, I just got a big kick out of that. You know, we've talked about the problem of security versus convenience. Typically we're worried about law enforcement forcing your thumb onto your phone, not your six year old. So parents beware.

Leo: So funny. So funny.

Steve: Apple announced last summer, at the 2016 Worldwide Developers Conference, that at the year end, that is, this year end, they would be requiring something known as ATS to be applied to all iOS apps. And they have learned - this is the story from the forcing people to change is difficult department, which of course is one of the interesting characteristics of IPv4 and SHA-1 certs and stuff we talk about all the time. It's like, okay, do we really have to change this? Apple's head of security engineering and architecture, Ivan Krstic, said during a WWDC presentation: "Today I'm proud to say that, at the end of 2016, App Transport Security" - that's the ATS, App Transport Security - "is becoming a requirement [fanfare] for App Store apps." He continues: "This is going to" - I added the fanfare. "This is going to provide a great deal of real security for our users and the communications that your apps have over the network."

So as our listeners probably already gathered, ATS is a feature which Apple debuted in iOS 9. When ATS is enabled, it forces an app to connect to web services over an HTTPS connection rather than HTTP. So it's very much like the HSTS we often talk about for web browsers. But of course it's forcing this behind the scenes. One of the problems, of course, is that we don't know what the apps are doing. And unless you capture packets and traffic and analyze it, the user is completely oblivious. At least with a browser and the URL we can get some sense for what's going on, although now that we've got JavaScript, it can be doing its own thing, too. But there's a lot of control over that.

So I wrote on my show notes: "Just as the 'S' in IoT stands for 'security,' as we know, the 'S' in HTTPS stands for 'secure.' But since we have little idea what mobile apps are doing behind the scenes, it can be impossible to determine whether an app's own cloud connections are authenticated and encrypted." So ATS is enabled by default for iOS 9, but developers have been able to switch ATS off and allow their apps to send data over an HTTP connection. That allowance was supposed to end at the end of 2016. And ATS requires the use of TLS v1.2, with a few exceptions for already encrypted bulk data, like media streaming.

But as the deadline approached, and as so many other people have found who've tried to set deadlines, Apple had to change their tune. On the 21st of December, Apple posted in their developer site: "App Transport Security, introduced in iOS 9 and OS X v10.11, improves user security and privacy by requiring apps to use secure network connections over HTTPS. At WWDC 2016 we announced that apps submitted to the App Store will be required to support ATS at the end of the year. To give you additional time to prepare" - because of course you'd only had six months at that point - "this deadline has been extended, and we will provide another update when a new deadline is confirmed."

So they haven't said when. They've just said, uh, okay, we're not going to do that. And no doubt there was a whole lot of backchannel screaming going on from people probably far and wide, saying no, no, no, no, no, we can't, we can't, we can't. Don't make us. And so Apple's like, okay, but, you know, we're serious about this. So get it done. So again, one more instance of just how difficult it is to get people to change the way they're used to operating.

So through the holidays there's been all of this talk about, as we were talking at the top of the show, Russian malware. I mean, this is a huge issue in politics at the moment with McCain rumbling about pursuing investigations; our new President-elect saying, oh, computers, nobody knows what's going on in computers because they're too confusing.

Leo: They're too fast.

Steve: His 10 year old spends...

Leo: Barron knows.

Steve: His 10-year-old son is a genius who could do anything with computers. But nobody really knows what's going on. But he does, for whatever reason, that's not getting any traction with him. So amid all of this, and I think because of this, a story got a ridiculous amount of attention, which was blaring headlines: "Russian Malware Detected in U.S. Electrical Utility." And in my notes I wrote, "Lots of smoke and noise." Some malware which has been associated with a known Russian hacking campaign, which has been labeled by the Department of Homeland Security and the FBI as Grizzly Steppe, S-T-E-P-P-E, was found on a single isolated laptop owned by a Burlington, Vermont electric utility. It wasn't in the grid. It wasn't crawling around. It wasn't doing a Stuxnet deal on us. They did an antimalware scan of all of their assets. And, oh, look, we've got malware on a laptop.

Well, okay, that's all it was. No indication that this thing was ever in any way associated with grid hacking or anything. It was just - it happened to be on a laptop that was owned by this power utility. And the laptop had and has nothing to do with the electric power grid operation and management. So everyone take a breath.

On the other hand, speaking of hyperventilation, Vermont's state governor, Peter Shumlin, said in a statement: "Vermonters and all Americans should be both alarmed and outraged that one of the world's leading thugs, Vladimir Putin, has been attempting to hack our electric grid, which we rely upon to support our quality of life, economy, health, and safety." And one of the state's U.S. representatives, Peter Welch, a Democrat for Vermont, said Russian hacking was "rampant, systemic, relentless, and predatory."

Leo: No. One laptop.

Steve: I know.

Leo: But that does comport with what we'd already heard, which is that - and what we'd expect, probes from time to time, just to look for weaknesses, to see where we could get some spearphishing done. I mean, nobody - I think we're going to rapidly enter into this age of mutually assured destruction, where we've got hackers; you've got hackers. You take down our grid; we take down your grid. So I think what you're going to see is exactly this kind of thing, which is, let's see. Let's nibble the edges. Nothing that would be provocative. Nothing that would be seen as an attack or an act of war. But let's see where the vulnerabilities lie. Does that seem sensible?

Steve: Yes. Well, so just to finish, Peter said: "They will hack everywhere, even Vermont."

Leo: Yes, especially Vermont.

Steve: Yeah, even Vermont has computers, "in pursuit of opportunities to disrupt our country."

Leo: Right, right.

Steve: And so in my notes I said, okay, great. So we've been talking about the inherent vulnerabilities of our U.S. power grid for, what, at least a decade. Yet no one appears to have the will to find or raise or commit the money that's going to be required to fix it. So, I mean, I'm glad for this, if this - I mean, I'm not happy that there was malware on the laptop. But if this, you know, if these politicians will please bottle their furor and...

Leo: Do something.

Steve: ...send it to Congress, and spend a lot of money on infrastructure that apparently will help the economy - that's what Trump says he's going to do - wonderful. Let's fix our grid. Let's not just keep talking about how incredibly vulnerable it is. And again, people get malware on their laptops, on their computers all the time. I agree with you, Leo, it is probably not a coincidence that this laptop had malware from Russia. But malware has to come from somewhere. And the fact is a lot of it comes from Russia, even if Putin had nothing to do with it, doesn't even know about it.

Leo: So you would kind of agree with Trump when he says, "Hey, this stuff, nobody really knows what's going on."

Steve: You've heard me saying now, it's been our theme for the last, at least half a year, the degree to which our systems are porous.

Leo: Right.

Steve: I mean, they just are. And they are too complex. We run across, like, you know, that bizarre way of getting into the iPhone we talked about, where you've got to touch your nose and click your heels three times and at the same time sneeze in a way that Siri recognizes, I mean, and oh, my, look, suddenly you're in. It's just because this stuff is just so complicated.

Leo: Well, and the challenge with the grid is that it's not one company. There's not one national grid. There's privately and publicly held companies all over the country.

Steve: And Leo, it's still got PDP-11s running it.

Leo: Right.

Steve: Which actually is a good thing. It turns out that only PDP...

Leo: Actually, that'd be safer, wouldn't it.

Steve: Yes. Only PDP-11s are qualified to control nuclear reactors because you can't infect them. So the real problem here is that it's old; you know? And we have an aging infrastructure. Our electric grid is, yes, a national asset. But it predates the Internet. And so what has happened is, after the Internet had happened, but before today, where a real appreciation of security exists, a lot of connectivity was created between the power grid and the Internet because, oh, look, we can check our meters from home, says the head of the electric utility for Vermont. Isn't that handy. I can make sure everything's fine on my iPhone. Just like you can check your baby monitor, and it can be hacked, too.

So we have a problem, and that is that we are still putting features ahead of security. And I expect it's going to take an attack on the power grid to bring about a change. I mean, I hope I'm wrong because widespread power failures are problems. And of course they could be made part of a larger terrorist plan or something in order to cripple response systems and so forth at the same time. So it really is bad. And I hope that, if nothing else, this story generates more concern. Our podcast listeners recognize that this is taking it a little too far, that this wasn't Russian Stuxnet-esque software that was found spinning on a turbine somewhere. It was on some guy's laptop that was in charge of the cafeteria. I mean, who knows?

Leo: Well, and it's reasonable to, I mean, anything that arouses alarm. The thing that I worry about is that the President-elect thinks that, well, that just means we should use handwritten notes and couriers as opposed to saying, well, let's try to modernize the infrastructure.

Steve: You know? And we want to avoid pigeons because apparently...

Leo: They're hackable, too.

Steve: ...pizza is now being delivered by drone, and that could be rough on the pigeon industry.

Leo: But you know what, this is what happens when you've had infrastructure longer than anybody else. You have piecemeal.

Steve: Right. That's exactly right. The reason Chinese airports are gorgeous is that they're all new.

Leo: They just built them.

Steve: Yeah, exactly. So I've been saying for some time that the concern over password length can be overwrought because really good, really, I mean, truly high-entropy passwords - which a user cannot create. You need software to do that. We just, you know, you've got to roll the dice, literally, or you've got to ask some software to give me some random gibberish because, if we do it, it'll be our Star Wars name by mistake, and we're in trouble. So I ran across a tweet from somebody who really knows his stuff. We've discussed Jeremi in the past. Jeremi Gosney is one of the, if not THE, leading password-cracking guys. He's affiliated, well, he's a principal and founder of two companies. And this first company, he calls it Sagitta, S-A-G-I-T-T-A, HPC.

So their description of themselves is: "Sagitta HPC is the leader in high-performance password cracking. We deliver enterprise-grade turnkey solutions that are designed by world-renowned password cracking experts and are tailored for information security, forensics, law enforcement, and litigation support professionals." So everybody knows how to read between the lines there; right?

"Our modular distributed solution can accommodate clusters of any size and integrates seamlessly with the popular free software you already know and love." Yeah, because these guys wrote it. "Whether you need a standalone system with three GPUs or a cluster of 300, you can count on Sagitta HPC to deliver the perfect solution.

"Sagitta HPC is a wholly-owned subsidiary of Stricture Group LLC, founded in January 2013 by Stricture Group founders Jeremi Gosney and Russell Graves, after a large number of inquiries were received asking them to replicate and improve upon their 25-GPU VirtualCL cluster. Since then, Sagitta has delivered solutions to dozens of government and law enforcement agencies, Fortune 500 companies, security consulting firms, and litigation support firms around the world." So what we're talking about is world-class, state-of-the-art, we don't care what it costs, we need to crack a password technology.

"At Sagitta's R&D lab," they write, "we perform heavy research into the best possible hardware and software combinations for our own internal use at Stricture Group. The best of the best solutions then become products that we make available to our customers. We push the bar higher and higher with each generation, frequently requiring us to write custom code such as od6config to enable the use of next-generation hardware. We also develop our own in-house code to maximize the performance and enhance the potential of our products. Sagitta also gives back to the community by frequently contributing and volunteering time to free and open source password-cracking projects, such as Hashcat and John the Ripper."

And I have in the show notes here a picture of Brutalis. Brutalis is their high-end, 3U-high rack monster. The description of this piece of equipment - and just looking at the picture I can see one, two, three, four, five, six, seven, eight double-width GPUs with lots of power cabling going back to a back plane; 3EEE redundant 1,000-watt each power supplies. The caption says: "Brutalis is an eight-GPU monster, clawing its way through hashes at unprecedented speeds. Providing up to eight Nvidia GPUs, two Intel Xeon E5-2600V3 CPUs, and up to [wait for it] 768GB of registered ECC memory" - so three quarters of a terabyte of RAM - "the Brutalis is the fastest, meanest, most hardcore system money can buy. Ships with a three-year warranty."

Leo: Now, Steve, you remember, somebody dropped by the predecessor to this.

Steve: Oh, right.

Leo: This is an incompleted board that was designed to crack DES.

Steve: Right. And I think only one of the chip spots was populated.

Leo: Yeah. There's one chip in here. But something - probably was a defect or whatever. But he brought this by, and this was the Cypherpunks out of Berkeley were making this DES-cracking machine. But I bet you this has one one-millionth of the power.

Steve: No, no. I'd drop the decimal point [crosstalk].

Leo: Yeah.

Steve: Oh.

Leo: The Brutalis.

Steve: This thing would just melt DES. It wouldn't get past the D. So with all of this background, now you know who Jeremi is. This is Jeremi. He tweeted: "I've encountered several people lately who use password managers and are generating random passwords, 20-plus characters long, some as long as 200." That's all he said. But the message there, the point is he, even this guy knows that, if your password has a large alphabet, meaning you can get upper and lower, special characters, and numerics, and a good source of entropy produced it, you don't have to worry if the website will only give you 16 characters. Because I see people complaining about 16-character passwords all the time.

Yes, if you're trying to fit your mother's maiden name and your first pet's name into 16 characters, you have a problem. But if a high-entropy source has generated that, even Brutalis will have a problem because the address space, I mean the attack surface, is so big on a 16-character - so it's, what, 95 is the typical alphabet size when you've got everything engaged. So that's 95^16. That is a very big number. And even a high rate of guessing, which Brutalis would bring to the party, is still going to make it very difficult to crack. And we hope then that the website did some PBKDF2 work on the backend. That is, it took your very, you know, your sufficiently long, very high-entropy password, and then made every single guess difficult. So anyway, I loved the picture of this monster machine.

Leo: I love the name.

Steve: It's like, oh, Brutalis. Wouldn't it be fun just to have one sitting over there thinking about something.

Leo: Yeah.

Steve: So Netgear's in trouble again. A security researcher, Pedro Ribeiro, discovered vulnerabilities, multiple unfortunately, in Netgear's WNR2000 routers, including a zeroday flaw that could be exploited remotely. Now, actually, as I read this now, the headlines I saw said zero-day. The text says that. But that's only, as we know, true if it's actually in use, that is, if it was found being exploited. So I think that's not what I mean to say. Let's not call it a zero-day flaw, but it is a present flaw. It can be exploited remotely to take full control of the device if remote administration is enabled.

Now, okay. No one should have remote administration enabled. What was somewhat surprising is that his scan found at least 10,000 vulnerable routers, that is, vulnerable Netgear WNR2000 routers, even though the current firmware, v5, has remote administration disabled by default. And in fact I have to think that remote admin has been disabled by default for some time. So unfortunately, these are not routers owned by listeners to this podcast because I know no one has turned on remote admin on their routers.

Leo: I hope not.

Steve: So what must be happening is that maybe there are some ISPs that deliberately configure them with remote admin so they can do remote support for their non-technically literate customers.

Leo: You bet. My Comcast router, cable modem router, my business-class cable modem router is writable from Comcast. And they insist on that so that they can, you know, "fix it" if I should break something.

Steve: And so that's one reason for the - because I'm a Cox cable subscriber, and it's the same situation. But of course I have a separate modem and then a little PC running pfSense. So they're welcome to fuss around with the cable modem all they want, but they're not going any further than the front door, you know, the WAN interface of the pfSense Firewall.

So there's the sad thing. He attempted to responsibly contact and notify Netgear of his findings, then decided to publish the advisory and release exploit code, which he has done, when Netgear never responded to his emails. So that's difficult to forgive on Netgear's part. The vulnerabilities that he found were in the Netgear WNR2000v5, which as I said before does not have remote admin enabled by default on its latest firmware. So users have to enable remote admin. Probably that's how this happened.

And what he discovered was that the web service, which in this case is uHTTPd is the daemon that's running in the router that provides web services, it is exposed to the WAN interface if remote admin is enabled. Probably you connect to port 80, and it shows you a web page where you log in. And so it turns out that you are able to get to the CGI scripts that are on that router behind the web server and leverage them to do all kinds of things

like change the password, reboot the router, and other things that should require authentication which don't. And he found a stack-based buffer overflow that he was able to leverage into a remote code exploit. So not good. He did a scan, found 10,000 of them with this front door wide open to the public Internet, told Netgear. They never replied. So after giving them time he said, okay, and he's now gone public with it.

So I don't know what you do in a case like this. I mean, this is so bad for this kind of user. Ten thousand routers are no doubt going to be taken over in short order because Netgear can't be bothered. I almost think in a situation like this a security researcher should just be quiet. That is, if Netgear's not going to respond, then he's done what he can to notify them, but it's going to take them to fix the problem. I guess you could argue that going public forces their hand and will get them to fix the problem, which I guess that's better. But, boy, you know, that one's a real toss-up because, yikes. I mean, basically the fault is Netgear's for being a provider of widespread networking routers and not being there to handle security.

I imagine at some point we're going to see some legislation. I mean, I don't know how we're going to get around some sort of requirements about security because we keep running across stories, for example, the Mirai botnet that brought down the East Coast DNS services, that was based on IoT and apparently some routers and cameras and things, and a DVR. Yikes.

Okay. So in the future we will do a deep dive into TLS v1.3. We've covered over the years the progression of SSL. We've done several podcasts about the secure sockets layer, SSL. TLS of course stands for Transport Layer Security. It has been at 1.2 for about eight years now. And the news is that TLS v1.3, the design was just finalized. The biggest practical development in crypto, argues the EFF, for 2016, that is all of last year, was Transport Layer Security v1.3. TLS, as we know, is the most important and widely used cryptographic protocol and is the backbone of secure Internet communication. I guess I would say that we know that it's built on top of TCP, which is built on top of IP, the Internet Protocol, which then carries the TCP, which then carries, once upon a time SSL, now TLS, in the multilevel or multilayer hierarchy.

So after years of work by hundreds of researchers and engineers, the new TLS design is considered final from a cryptography standpoint. The protocol is now supported and available in Firefox, Chrome, and Opera. And although the naming, you know, calling it TLS 1.3 makes it seem like a minor version upgrade, it is, and I'm really happy for this, a major redesign from TLS 1.2, which, as I said, is about 10 years old. In fact, one of the most contentious issues was whether the name should be changed again, much as we changed SSL to TLS, specifically to indicate how much an improvement TLS 1.3 really is.

On the user-facing side, we will probably see, once support is ubiquitous, a noticeable improvement in speed. TLS 1.3 has been tuned for speed by incorporating a lot of earlier research, reducing the number of network packet roundtrips required before data can be sent, either down to a single roundtrip or, in the case of repeated connections, zero roundtrips. And these ideas, as I mentioned, have appeared before in experimental form with the QUIC protocol. And there was another one called False Start for earlier versions.

So those were sort of grafted add-on, experimental add-ons to earlier TLS versions, but now have become part of the default behavior of TLS 1.3. And they'll become, over time, much more widespread. And of course this is important because our web pages are getting heavy. They're increasingly consisting of a huge number of individual assets, each of which requires - increasingly is coming from different locations and so requires different connections to be set up. So this should reduce latency and improve web page loading times.

And there's also big improvements in security. It incorporates two important lessons, which is music for me, from decades of experience with TLS. First, the protocol has been simplified by removing support for a number of old protocol features - yay - and obsolete cryptographic algorithms. And it was also designed with the benefit of a developmental technology known as model checking, which has been used to find flaws in many older versions of TLS and SSL. But TLS 1.3 was analyzed extensively, using this model checking, by the cryptographic community before the standardization process, instead of, as has been done until now, waiting until the protocol is widely deployed and thus, as Apple as found, difficult to patch in the field.

So this is great. You know, it is, I would argue, this is the core protocol of the Internet, and even more so going forward. As we know, it provides authentication and privacy, those two things we have to have together in order to know who we're talking to and in order to protect the secrecy, the privacy of whatever information is exchanged. So in the future we will do a detailed walkthrough about what the new features are in 1.3. I just wanted to mention that it exists, and it's coming soon to all of the technology that surrounds us.

Leo: Nice. We've got some teasing, exciting things to come up, including a show that you said I have to watch.

Steve: Yeah, and I'm batting a thousand so far in recommending it. So SSD, easy to remember, ssd.eff.org. That's the domain name that they gave their page, ssd.eff.org. I can't recommend this highly enough. SSD is the abbreviation for Surveillance Self-Defense. It's in three broad categories, to give you a sense for what's there. Overviews is a section that contains subtopics: An Introduction to Threat Modeling; an Animated Overview: How Strong Encryption Can Help Avoid Online Surveillance. Another one: How to Make a Super-Secure Password Using Dice.

Leo: I love it.

Steve: Yup. Another animated one: Protecting Your Device From Hackers; and Using Password Managers to Stay Safe Online. Then there's, also in Overviews: Choosing Your Tools, Creating Strong Passwords, Keeping Your Data Safe, Seven Steps To Digital Security, What Is Encryption?, and Why Metadata Matters. Then they have a whole series of how-to tutorials: How to Avoid Phishing Attacks; Circumvent Online Censorship; Delete Your Data Securely on Linux, on Mac OS X, on Windows; Enable Two-Factor Authentication; Encrypt Your iPhone; Install and Use ChatSecure; Use KeepassX; Use Off the Record (OTR) for Mac, for Windows, for Linux; Use PGP for Linux, Mac, Windows; Use Signal for Android, Signal on iOS; Tor for Windows, Tor on Mac; WhatsApp on Android, WhatsApp on iOS. So really comprehensive. And then they have what they call Briefings for the third category: An Introduction to Public Key Cryptography and PGP; Attending Protests (International).

Leo: Oh, I like that, yeah.

Steve: Yeah. Attending Protests in the U.S.; Choosing the VPN That's Right for You; Communicating with Others, wonder what that's about; How Do I Protect Myself Against Malware?; Key Verification; Protecting Yourself on Social Networks; The Problem with

Mobile Phones; and Things to Consider When Crossing the U.S. Border. So just a really, really neat-looking, comprehensive, across-the-board set of coverage of interesting things from the EFF.

Leo: Now, I notice this video comes from Al Jazeera Plus. So it sounds like they've collated material from a variety of sources.

Steve: Oh, right, right.

Leo: As opposed to all original. But I'm sure it's good. These guys know what they're doing.

Steve: Yeah.

Leo: Yeah. Really, really nice stuff. This is ssd.eff.org. And if you use it, and you tell your friends, you might consider becoming a donor, as I am, to EFF. I think they do really good work, and they're one of my monthly charities. And I feel good about it when I see stuff like this.

Steve: Yeah. So as I said, next week our topic will be a look into this Russian PHP-based malware. We'll share what we know, and then we'll appraise what we think of what we know.

Leo: Yeah, that's what I'd like is, I mean, Dan Goodin...

Steve: To put it into context.

Leo: Yeah, in Ars Technica basically mocked the report, but also pointed out that it's possible they couldn't actually give us a substantive report without revealing techniques and so forth. So I don't know.

Steve: Yes. And this is the problem.

Leo: It seems like a very sloppy report, to be honest with you.

Steve: Yeah. Well, it looks like the breakdown of the malware was really itself interesting.

Leo: Yeah.

Steve: So I think there's probably deliberate obfuscation. And the problem is we just, on

the outside of the inside security community, we don't know how much we don't know. We don't know what is known. And as we often talk, attribution is difficult. Well, for example, the fact that malware associated with Russia was on that laptop doesn't even mean that it came from Russia.

Leo: It's from Russia, yeah, of course not.

Steve: I mean, it just, you know. And as I said, a lot of malware is Russian. A lot is Chinese. I assume that there's a lot that the U.S. is doing, too. I mean, say Stuxnet, for example. I think that qualifies. If you're on the other side looking at us, that's malware from the U.S. So, yeah. Anyway, the technology is what's interesting. The rest is just unknowable. And there's really not much we can say one way or the other. And we know how difficult attribution is.

So I do have some fun miscellany. And first of all, last week's rerun from years before of the Portable Dog Killer episode was a huge success. Many people enjoyed hearing it again because it had been so many years since they had. And, as I expected, since then we have acquired many more listeners, for whom it was the first time through. And I got a lot of great feedback from that. One question I saw many times caught me up short, and I thought, isn't that interesting. And that is, you don't have a picture of the Portable Dog Killer? And I bring this up because I realized, no. That was 1970.

Leo: In those days you had to have a camera with film in it.

Steve: Yeah. And photographers had cameras. And pretty much there was - you might have a Brownie.

Leo: Instamatic.

Steve: Or an Instamatic, yeah. And I thought, isn't that an interesting change, that now no one doesn't have a camera. And photographers were like a profession. I mean, they still are, obviously. But you didn't have just everybody with a camera. And so, yeah, I mean, I've got pictures of what I eat every day, but I didn't have a picture of the Portable Dog Killer because no one took pictures of everything. I mean, first of all, there was nothing to do with it. If I had a picture, I could show four people. But there was no Internet. And so there was nothing to do with photos. There was nowhere to put it that it would be seen. So I thought, wow, isn't that interesting that, I mean, I know in my head I can still see it. I know exactly what it looked like. I described it during the podcast. I still see it clearly. But it didn't even occur to me to take a picture because why? I had it. And there was no Internet that would allow any kind of a broadcast of a picture.

Leo: Right. Even if you had a picture, it's in a shoebox somewhere in Mom's house.

Steve: Yeah, exactly. I mean, if I'd been on the news or something, then it would have been broadcast. But that was literally, when you think about it, look at the change now where everyone has access to broadcast technology. And back then there was none.

Leo: Yeah, it's amazing. No trivia unshared these days.

Steve: Yeah.

Leo: Hey, are you in chat?

Steve: Yeah, I have it running in the background.

Leo: Okay, good. So people are talking to you in chat. And I said, "Oh, that's not Steve. He's never in chat during the show." But then I looked at who is, and it's somebody coming from Cox in Orange County, so that makes sense. All right.

Steve: Yes. And I fired it up, and then I put it behind my PDF here. It doesn't know me, so I can't ever say something quickly, so I just left it running because...

Leo: We can handle that. In fact, what you probably - I'll defer to the chat mods because I don't really run this thing. I don't know how it works. But you probably want to register that name so that no one else can use it. And at that point, once you register the name, the mods can make you permanently voiced and opped.

Steve: Oh, okay, cool.

Leo: Yeah. But you'll need to investigate that.

Steve: Okay. Looking back on 2016, I also wanted to remind people that, without a doubt, The Sequence was the Puzzle of the Year. As we know, puzzles have become a staple of the podcast. And I get lots of referrals. I look at them. Most of them don't stand up. But The Sequence, not only did it stand up, no puzzle has generated anywhere near that much positive feedback and satisfaction and happiness and joy. I sent gift links to my brother-in-law and nephew, his son, after Christmas, just because I thought, you know, I'm sure they don't know about it. I want them to have copies. You know, it's 99 cents.

So if anybody missed that, if they're listening to the podcast since then, or they didn't get around to it, I just thought I'd say, you know, it is wonderful. It is essentially sort of - it's visual programming. You're solving an animated visual puzzle with little widgets that each do their own thing, and you have to put them together in the right place to move the little hockey puck around. And it's just - it's a perfect puzzle. So I just wanted to acknowledge that The Sequence was the puzzle of 2016. And without question, the most retweeted thing that has ever been said, and definitely the top slogan of 2016, was "The 'S' in IoT stands for 'security.'" Everybody loves it.

Leo: And I did not get it at first, so there you go.

Steve: Yeah, well. So thus its subtlety. Okay. On December 23rd, two days before Christmas, Netflix dropped a new series. I was up till 3:30 a.m. Monday morning because I couldn't stop. It was just, okay, just one more. Okay, okay, just one more. "Travelers." I didn't know why it was so good. But the name Brad Wright that is prominent on the screen, he produced the show, and he's the writer on many of the episodes. I thought, why is that name so familiar to me? Well, he's a Canadian television producer, screenwriter, and actor, best known as the creator and co-creator of the television series "Stargate SG-1," "Stargate Atlantis," and "Stargate Universe."

And I've mentioned "Stargate" before. It is one of the best, little-known, often-missed, really well-produced and assembled science fiction series. Whereas Kirk used the Enterprise to fly around and get in trouble, these guys used a discovered alien artifact which allowed wormholes to jump them around. And so it was very much the same sort of vehicle. You went somewhere. You dialed a gate address, and you didn't know what you were going to find. So they'd send a little probe bot through to see if the air was breathable and look around; and, if so, the team would go through. Anyway, fabulous series. So I thought, okay, that's why I know his name. And that's why "Travelers" on Netflix is amazing.

Leo: Sci-fi?

Steve: Sci-fi. Now, when I first heard the description, and this is not a spoiler, you know, I don't do those except maybe for this week's picture, unfortunately. But I just had to, it was so good. But so the non-spoiler background is, in the future, we have really messed up things. And while...

Leo: The future?

Steve: In the future.

Leo: Why wait till then?

Steve: What could possibly go wrong?

Leo: The future? Okay.

Steve: So, like, most of humanity has been killed off, for example.

Leo: I'm hoping we'll have unscrewed up things by the future.

Steve: It doesn't look like that's where we're headed. So, though they cannot send things back, because of course you can't, it turns out they can send personalities back. They can send people packages, whatever you want to call it, to essentially take over people in the past. And so when I hear that, I go, oh, no, what a cheesy concept. I mean, like how could you come up with a less expensive way of doing science fiction

than suddenly have an actor go, "Oh, I'm different now. Now I'm from the future." It's like, okay.

So if you watch the first one, be warned. You may not sleep. There are 12 episodes in the first season. Yes, there will already be a second one, but we're going to have to wait a year. Pace yourself if you can. I was unable to. I managed to cut myself off after four hours on Sunday night. And then when I dipped in on Monday I could not stop. So I went till 3:30 a.m. and finished eight episodes the second night. It is surprisingly well written. It's the writing. It incrementally reveals more information. What I told you, you figure out almost immediately, so thus it's not a spoiler.

But believe me, I mean, we loved "Stranger Things." This thing is right up there. I told Jenny, Jenny and her mom. I got a text from her this morning saying they loved the first hour. Three other people I've told just immediately went crazy. I tweeted it two nights ago at the beginning of Night Two for me. I said, look, I'm going to be talking about this on the podcast Tuesday, but you may not want to wait. I got a lot of people who agreed with me. One person said it started out slow. And I'm thinking, okay, what show are you watching? Maybe he had the wrong - "travellers" with two L's will get you something.

Leo: He was hoping Arnold Schwarzenegger would be in it.

Steve: I don't know what's going on. But anyway - oh, and it's also produced by and stars Eric McCormack, who is, of course, Will from "Will and Grace."

Leo: Oh, I love him. He's great.

Steve: So, yeah. He's a great actor. The acting is good. Just I can't tell you, just it is so good. So unreserved recommendation for "Travelers" on Netflix.

I did want to give everybody a heads-up that in two weeks we get, finally, in fact, it had been off my radar for so long I'd forgotten about it, Season 6 of "Homeland," Carrie coming back. Because "Homeland" is just a great series on Showtime. And "The Expanse," the second season of "The Expanse," which was a very well-liked sci-fi series. I read the whole book trilogy, and then they kept on writing more. The first season we were all remarking was really well done, surprisingly well done. I was very impressed with it. And it's a cool concept that we were just about to find out about when the first season ended. So you have one month, if you want to rewatch the first season to get ready, because Season 2 begins on February 1st, so four weeks from today.

And SpinRite. I didn't get a chance to talk about this. This was in my notes from last week. Yet another thing that in all the times we've talked about SpinRite, I have never mentioned. It was introduced on December 18th by someone, Jason, who tweeted me a picture of something he says he has never seen before. Actually, a friend of his. So he said: "@SGgrc Friend of mine says he's never seen cabling errors before. What's going on here?" And then he sent me a picture, which is in the show notes, which shows very healthy-looking SMART data with ECC corrected and relocated sectors all up at max. But SpinRite is detecting cabling errors and showing the total of - there were more than 19,000 of them at however far that had gone, with a pretty constant rate of them.

And that's something I've never talked about, but it's another one of the cool things that SpinRite will do because, think about it, and it's something we don't tend to think about,

you need a good cable between your drive and your motherboard. You've got smarts in the drive now; but you need to make sure that, when you're sending 4,096 bits up the cable to create 512 bytes for a sector, that what the motherboard intends to send, the drive receives correctly. And if you had a flaky cable, you could be recording data incorrectly, even though the drive was recording what it received correctly. So since that was available, of course I brought it out to the UI.

And so what happens is there is, in the protocol, an error correction. I'm sorry, it's not error correction. It is essentially a CRC, a cyclic redundancy check, which the sender appends to the data, and the drive receiving it verifies. And if it gets a CRC fail, meaning that the data that the motherboard sent had an error in transit, then the drive, thank goodness, will spot that and fail the transfer. And inside it, it counts, in its own data, that it received the error. So when SpinRite queries the SMART data for all of these variables, if it sees that the cabling error count is increasing, it immediately locks onto it and then starts tracking it and showing you the number and rate at which those are occurring as SpinRite runs.

So it's another example of something where it only shows up under use, in the same way that error correction or ECC retries or the dynamic relocation of sectors only occurs when it's in use. Similarly, the cabling errors, you're not going to see it if you just, like, look statically. You need to put it under load. So just one more benefit of SpinRite, and something that it's really fun to have surface out on the UI.

Leo: The topic of the hour really was started some time ago when somebody emailed us, right, about those weird CAPTCHAs, where you click a button and say "I'm not a robot."

Steve: So, yes. I was using the 'Net and encountering this checkbox. And I'm thinking, wait a minute.

Leo: Are they secure?

Steve: It just asked me if I'm a robot?

Leo: No, no robots here.

Steve: And I just say no? And every time I click it it goes, okay, and then lets me do what I'm going to do because after all, despite what some people may say, I'm not a robot. And so I'm thinking, okay, what's going on here? Because, I mean, we've covered CAPTCHAs. We did an episode on it, this whole idea of coming up with a way to prove that you're a human because many situations don't want automated bots to do stuff. And of course we were talking last week, or, sorry, two weeks ago, the last podcast, about the ticket purchasing business and how finally legislators had said, okay, we're just going to make it illegal for bots to bypass anything that is trying to prevent a bot from purchasing so that, if the seller of a ticketing website wants to only sell to people, and they implement any technology to distinguish people from robots, it is now illegal to bypass that. So how does this CAPTCHA thing work?

Okay, well, first of all, this is a Google property. So this is a service, this CAPTCHA,

reCAPTCHA, I'm not a robot thing, a service from Google. They're not telling us exactly how they're doing it.

Leo: Oh, interesting.

Steve: Because they don't want, you know, the more you know, yes, this is a little bit of obscurity, but the more you know, the more possible it becomes to defeat it. So thanks to the fact that it's a server-side thing, there's no way to analyze what's going on on the web page and figure out what's happening. But there's been some probing done because of course we can probe it. So one investigator believes that using incognito mode blocks the easy checkbox. And I think that's probably true because I'll tell you why. Another investigator has partially spoofed the system by using a B-spline mouse path with randomized way points and destination. And it's been determined that the user's browser must be able to render the canvas, meaning that you must be able to use JavaScript to draw onto the surface of the web page. So, and those are like the bullet points from several hours' worth of digging that I did. And here's what I came up with. What Google is doing is everything.

Leo: Of course.

Steve: Yes, of course. Think about what Google actually knows about us. If this is a "I am a robot" thing, CAPTCHA, that is being displayed on my web page from a Google property, then that server got my Google cookies. So it knows who I am. It's got my Gmail. It knows the IP from which I'm making the request. It knows where I've been, what devices I'm using. Think about, like, what Google knows.

So essentially this is comprehensive reputation verification. And they are also monitoring the mouse pointer and looking at the rate of movement and the path the mouse makes to see if it looks automated, or does it look human-ish? Do you always stop dead center? Or as you click from time to time are you stopping in different places? So essentially they are doing a highly comprehensive analysis of your interaction at the moment, and your entire history that Google has on tap based on your Google identity, which is sent from your browser because this property comes from Google. So your browser sends a cookie. Google says, oh, you know, Steve logged into Google Drive to prepare the show notes an hour ago, and he's been apparently homebound over the holidays working on stuff, and this is his IP that hasn't changed since the power outage three months ago when his cable modem came up with a new IP. So, yeah, we're really pretty sure this is not a robot.

So with a final verification of watching the mouse move and him clicking on the box, we will draw some stuff there, make sure that the surface he's viewing is renderable using the canvas API, all that goes well, yup. And look at what we get in return for that. We just get a checkbox. No more house numbers at an oblique angle at midnight with a poorly lit light. And, I mean, there are some CAPTCHAs, you look at them and go, I'm really sure I'm not a robot, and I still have no clue what that says. So that's the story. And once you see it, it's like, oh, yes, this is what you would want to do.

Only someone like Google could, that is, or Facebook could. But somebody who has a comprehensive background based on behavior and IP address and constant interactions with their service that occur frequently enough that they're able to develop a comprehensive belief in who you are. And I'm sure, if you started suddenly to behave

like a robot, they'd mark that down in your "I'm not a robot" skepticism column in some database somewhere, and you'd have to be filling out - oh, and by the way, if it's not sure, you still will be prompted. If you click "I'm not a robot," and it goes, eh, okay, after this question, you can still get a picture and have to fill out what the numbers are, if you need to prove to it.

And so, for example, this is why incognito mode blocks the checkbox. Incognito mode strips your queries of cookies, so your query is coming in, maybe from a known IP, but that's not enough. Google needs to know the logged-in user. And so it'll say, you know, we want to have a reliable "I'm not a robot" detector. So if you use incognito mode, you've got to fill out this little questionnaire here. Otherwise, just check the checkbox. So I thought that was cool.

And again, it's like any, as we originally talked about it, this is a difficult problem to solve if you constrain your solution to just the instantaneous interaction of the user and the web page. And we could argue it's probably impossible. It's an intractable problem because, for example, we've seen these things being exported to other countries in sweatshops. Then all they do is spend all day typing in what the image says and exporting the results back. So it's spoofable. But this is not spoofable. This is your full reputation and past history brought to say, yeah, you know me. I'm not a robot.

Leo: Well, it's funny because every once in a while I'll click the button, and it will ask me for more information. So maybe there are things sometimes where, I think, unless my memory's wrong, but I think I've clicked it, and it's occasionally said, okay, now you have a real CAPTCHA. So maybe I didn't satisfy its...

Steve: Yeah. And no doubt there's heuristics going on. Maybe your IP address had just changed. Maybe you were using, you know, you're constantly setting up new computers. So a fresh computer wouldn't have Google properties and Google cookies and things registered to it yet. So it could be lots of stuff.

Leo: Yeah.

Steve: And it's probably, they're probably sucking in your browser headers and saying, oh, yeah, we've seen that browser before. That's definitely Steve.

Leo: Yeah. It just points up how much Google does know about you.

Steve: Yes.

Leo: Yeah, yeah. Very nice, Mr. G. Good way to start 2017. And end right on time, I might note.

Steve: Uh-huh.

Leo: You are a master. I don't know how he did that. Back-timed it. Steve Gibson's at GRC.com. That's his website, the Gibson Research Corporation. That's where you'll find SpinRite, the world's best hard drive maintenance and recovery utility. You also find all his great cool stuff that he offers. SpinRite's the only paid thing. Everything else is free. And podcasts. All of these, all 593 of them, audio as well as transcripts at GRC.com, as well as the show notes.

We have audio and video at our website, TWiT.tv/sn. We are also on YouTube. You can watch us now on YouTube Live - thank you, YouTube - YouTube.com/twit. So you can watch live when we do the show, there or on our website, we have many live streams, or on one of those great apps. There's, like, five Apple TV apps. There's a Roku app. A bunch of apps.

You should tune in if you want to see it, though, every Tuesday, 1:30 Pacific, 4:30 Eastern, 21:30 UTC. I give you the UTC so you can do your own math based on your location. Join us in the chatroom, too, where Steve has unaccountably appeared, irc.twit.tv. I said, "Oh, no, that's not Steve. Steve's never in there." They said, "No, it is Steve." Register your nick. You can also, let's see, did I say everything? I think I did. Subscribe, that way you don't miss an episode, to Security Now!. Thank you, Steve.

Steve: I will see you next week to talk about what we know of Russian malware.

Leo: I can't wait. And I'll be watching that show all week long.

Steve: Oh, boy. "Travelers." "Travelers." "Travelers." "Travelers." Yes. And we will do a - we'll talk about it next week.

Leo: Canceling my weekend plans now.

Steve: Perfect.

Leo: Thanks, Steve. We'll see you next time.

Steve: Okay, buddy. Thanks.