SECURITY NOW!

Transcript of Episode #503

Listener Feedback #210

Description: Mike and I discuss the week's major security events and discuss questions and comments from listeners of previous episodes. We tie up loose ends, explore a wide range of topics that are too small to fill their own episode, clarify any confusion from previous installments, and present real world application notes for any of the security technologies and issues we have previously discussed.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-503.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-503-lq.mp3

SHOW TEASE: It's time for Security Now! with Steve Gibson. The EFF wins its podcast patent challenge. Steve will fill you in on a massive botnet takedown and update you on the Mac Rootpipe vulnerability and so much more. Stick around. Security Now! is next.

MIKE ELGAN: This is Security Now! with Steve Gibson, Episode 503, recorded Tuesday, April 14th, 2015: Your questions, Steve's answers, #210.

It's time for Security Now! with Steve Gibson, the show about privacy, security, coffee and more. Steve, how the heck are you?

Steve Gibson: Very good, Mike. Everyone can tell from the voice that just introduced the show that we have Mike Elgan today co-hosting with me. Leo is at the National Association of Broadcasters Conference, or Convention. I guess he was a speaker over the weekend. And so he's probably hobnobbing and on his way back. But so Mike is filling in. We have a Q&A this week, our 210th Q&A for Episode 503 of Security Now!. And so we're going to hear from our listeners. I found 10 great questions, which Mike will read and I will then answer at the second half of the show. And we're going to cover the week's news. I sent a note to Leo and a whole bunch of email addresses of the TWiT people who I normally send show notes to every week over the weekend because the EFF, our friends at the Electronic...

MIKE: Frontier Foundation.

Steve: Frontier Foundation. I can never remember that. It just seems like an awkward acronym for some reason.

MIKE: They'll always be EFF to me.

Steve: Yeah, well, so EFF. Everyone's been a little bit on pins and needles over the podcast patent suits that have been launched by Personal Audio, I think was the name of the company, that is your classic patent troll. So we'll talk a little bit about the news over the weekend. We'll follow up a little bit on the CNNIC catastrophe with them issuing an intermediate certificate that was found in the wild to be itself issuing certificates, and how Google and others have reacted. We now have some probably final results from the other browsers aside from Chrome.

There was also news of a massive botnet takedown that's sort of interesting. A problem that has been fixed in the most recent version of Yosemite, Mac OS X, but apparently is not going to be fixed in older ones, known as the "Rootpipe" vulnerability. Kaspersky, working with some authorities, got a hold of a whole bunch of decryption keys for some ransomware that we'll talk about. The government, the U.S. government, is still trying to figure out how to negotiate with Silicon Valley and techno companies to find a way for law enforcement to be able to decrypt communications because they're just not happy with the idea that we could be talking, and they can't be listening.

And then I want to also follow up on what Leo and I sort of discovered on the fly last week about the disturbing ordering of cipher suites. It was BofA that we looked at where the BofA server was, like, offering the worst possible ordering of cipher suites, and a lot of our listeners have followed up on that, too. So a bunch of great news for the week, and then Q&A.

MIKE: I don't know about you, Steve, but I can't wait. I'm looking at some of this stuff, and it's so fascinating. And we cover a lot of this stuff on Tech News Today, and the difference is that you provide answers. On Tech News Today, we just scare the daylights out of people. So Steve, so, now, we're going to go into this list of news. Now, this podcast patent issue is one that's near and dear to all of our hearts because the idea that a patent troll could actually start mucking around with podcasting kind of gave us all the creeps. But what's going on with this?

Steve: Well, so without getting into all the details because, I mean, this - I attempted to read the legal mumbo jumbo, and it just makes your eyes glaze over because they're literally parsing sentences for meaning, and you really need to have a deep background. So I'm sure that Denise (TWIL) could probably take this apart a lot more. What I did manage to figure out is that what was demonstrated essentially by just having the EFF and their intellectual property-savvy attorneys prepare a statement to the U.S. patent and trademark office, is that they found two existing prior art, as it's called, patents, which essentially rendered the patent that was being used by this Personal Audio LLC to be ruled invalid.

So, I mean, this is exactly what we expected when this first surfaced was that here was a company saying, hey, we invented podcasting. And everyone's intuition was no, you didn't. Certainly somebody has been doing something like podcasting for far longer than the date of this invention. But it's one thing for us all in the industry to have an intuition about that, and another thing to be able to demonstrate it to people not involved in the industry, who are at the Patent and Trademark Office, who, if they don't see something clearly demonstrative that there isn't prior art, they'll just say, okay, well, yeah, this sort of seems new. We'll say yes, and we'll let the courts decide.

And unfortunately, this has been the problem is that, for those of us who do "engineer" software, and I use that term as opposed to "invent" because most of the time that's what this is. It is, if somebody trained in the practice is asked to solve a problem, we apply the things we know to solving the problem. That's engineering. And there are certainly moments of invention. There are things that represent inventive inspiration. I think maybe of the zipper or Velcro or something where, you know, you could be struggling with it for a long time and never come up with that. And then someone shows it to you, and it's like, oh. That's an invention. Instead of something where 10 people are asked to solve a problem, and they just all solve the problem. Which is not invention, it's engineering.

So anyway, the bottom line is that there has been a lot of angst surrounding this because these guys sued the three major broadcasting networks, and I think it was Adam Carolla was also named in the suit, saying that all of these entities were in violation of this patent troll's patent, and they wanted to collect licensing fees for podcasting. And of course the danger was, if they had been able to prevail in those lawsuits, then everybody else, all other podcasters, then there would be precedent for this patent being upheld. And it's easier in that case, rather than fighting a patent which has already been upheld, to probably capitulate and pay. And no one wanted to do that if, in fact, these guys didn't have a true demonstrable intellectual property right for podcasting.

And there were two existing, I mean, basically it was shown this was going on already at the time that this other group, this patent that these guys had, was supposedly invented. So they didn't invent anything. Or if they did, it was already in the public domain. And even if they didn't know about, the fact that it was known disqualifies them as having invented it.

MIKE: Now, a quick note about patent trolls in general and how this works in the United States. It turns out that the Eastern District of Texas is very - the judges and juries in that part of Texas are very, very friendly toward patent trolls for some reason. And so what happens is that companies like Personal Audio LLC, they don't really have a company, they don't have a bank of engineers inventing things. That's why they're patent trolls. They just, they basically buy patents and then they sue.

Steve: Litigate.

MIKE: Exactly. So they get a P.O. Box in Eastern Texas and say that they're headquartered there. Then this is where the lawsuits take place. And then they oftentimes win lawsuits that they shouldn't win. It's really an interesting, to me it's an interesting part of the dirty tricks that patent trolls actually use.

Steve: Yeah. I saw a little blurb about that. I thought it was funny, Samsung has been brought to trial there so many times and had such bad problems that they started to get much more involved in the community. They're, like, major financiers of, like, the local stadium, and they've got Samsung brings happy times, you know, like logo stuff all over the place, basically trying to turn the actual citizenry in that area to be a little more pro Samsung so they'll lighten up on all of these bogus decisions that are being fought against.

MIKE: That's hilarious. That is hilarious.

Steve: Yeah.

MIKE: In a bad way.

Steve: So, oh, yeah, in an unfortunate way. So we know from having covered this a couple times that at the beginning of, I guess it was about three weeks ago that alarm bells went off at Google when a fraudulent Google certificate was sniffed by someone's version of Chrome. That immediately started an investigation where they found that a major certificate authority in China, in fact CN, as in the top-level domain abbreviation of

China, CNNIC, had issued an intermediate certificate which allowed another company to sign certs on behalf of any domain they wanted to.

Well, that's an absolute breach of the trust that we put in a root certificate authority like GoDaddy, like VeriSign, or like CNNIC. So Google's response, after supposedly performing some investigation, talking to the CNNIC people, and we covered this last week where Adam Langley updated his blog, I think it was actually April Fools Day, and he updated his April Fools Day blog posting where he said - I don't think he even used the word "agreed" because that certainly wasn't the way it happened. But after a joint investigation, we've decided that we're no longer going to honor CNNIC's certificates. We'll arrange some sort of a whitelist for existing ones, but basically we're putting them out of business. The major browser on the Internet, which is now Chrome, will no longer allow secure connections to CNNIC customers.

Now, it may be possible to override those in some cases, which of course is a bad thing. You don't want to train people to ignore security certificate warnings. But the HSTS technology, part of that forbids the user from overriding. So in many cases you just won't be able to get to those sites. And here's the key. If they have a certificate issued after April 1st. So apparently they're looking at the signing date, that is, Chrome is, April Fools Day of 2015. And any CNNIC certs signed after April 1st will not be honored by Chrome. So that bomb dropped.

And we did find a comment, very unhappy, that CNNIC posted in their own blog, saying that it was incomprehensible to them that this is what Google was doing. I mean, it may have been a joint investigation, but it was certainly not a joint announcement. CNNIC was, from all appearances, furious with this. And you can imagine why. I mean, basically, they are out of business. No one is going to buy a certificate from them now because the majority browser on the Internet will not honor it. So everyone is going to get any subsequent certificates somewhere else. And so essentially this grandfathers in, by using the signing date, this grandfathers in the existing certs, but just says, you know, I mean, it's more than a slap on the hand. This is, sorry, you're out of business.

MIKE: Of course, Google said that they welcome CNNIC to reapply for trusted status, quote, "once suitable technical and procedural controls are in place," unquote. And the other thing that's interesting about this, Steve, to me is that this is kind of the quasi-official Internet authority in China. This is like, you know, this is...

Steve: Big.

MIKE: I don't know what's comparable there within the United States. But this is, you know, these guys are completely in bed with the Chinese government.

Steve: Yeah. Yeah. And the other thing we don't know from even the closest reading of what Adam posted, and I haven't seen anything in any greater depth anywhere else, you get the sense that what we're seeing is sort of the political frosting on top. And something must have gone on. And Leo and I were speculating on this without any facts because there aren't any, like what's the real story? What actually - what don't we know that underlies this decision?

Now, here's what's interesting, and the reason I'm bringing this up yet again, is that we now are three weeks after the fact. Apple has updated their root CAs in the most recent updates of the Mac OS and iOS. And CNNIC's root is still there. So Apple has not followed suit with Google, and it turns out neither has Microsoft. Microsoft has taken to doing something that - actually we have a question in the Q&A today that I put in there before I ran across this update, where essentially Microsoft delivers certificates on demand. And a fresh install of Windows 10 just yesterday was able to obtain a CNNIC cert from Microsoft on the fly. So Windows, Microsoft's products are also still honoring CNNIC certs. But Mozilla did exactly follow Chrome's behavior. And as of Mozilla's most recent update, they, too, will no longer honor certificates signed by CNNIC after April 1st.

MIKE: Now, it has to be said, Steve, that, you know, we follow lots of different stories that involve American companies and how they interact with China and the Chinese government. And this general trend line follows what has happened historically. So, for example, Microsoft really kisses the Chinese government's behind. Apple kisses their behind more than any other company in the world. They were under attack, Apple is under sustained attack from the official media in China over numerous issues that were all trumped-up, baloney charges. They accused Apple of discriminating against Chinese customers. They accused Apple of creating a security nightmare.

In every single case, Apple said, essentially, we sincerely apologize; we'll do anything you want us to do; whatever you say, we'll do it, because of course China is about to become Apple's biggest market. This is one of the reasons why Apple is the most valuable company in the world is China. And Google, on the other hand, has a history where they were even probably hacked by the Chinese government, and they left and moved their offices to Hong Kong. So Apple is one of the few Western companies that has fought the Chinese government, and here is yet another case. Although, again, this isn't the Chinese government, per se. But they're not willing to do as much as a lot of other companies are in terms of bending over backwards to please the authorities, whoever they are, in China.

Steve: Yeah, and, of course, although Google is certainly a commercial entity, Apple and Microsoft are, I mean, their definition is we're supplying operating systems and solutions that need to operate globally. And so, for whatever reason, they've just decided, hey, you know, a mistake was made.

I did also learn something that I never mentioned before, because I didn't know it, and that is that this intermediate certificate only had a two-week life. So it was some sort of a brief testing certificate. It's not like it was going to ever be alive for three years and all kinds of skullduggery could have gone on if no one had noticed. It was due to expire. It was issued for two weeks. And so it was just during that brief two-week window that it got installed in this proxy that allowed it to sign a Google domain as part of its proxying an HTTPS connection that set off the fire alarms and the alarm bells. So it even wasn't as bad as it could have been because CNNIC deliberately, under whatever terms and conditions, issued a two-week-long cert that wasn't going to be able to do anything after that time.

So anyway, again, my sense is there's, as you say, a lot of politics going on sort of behind the scenes. And Mozilla, I guess, arguably the least commercial of the four, just says, you know, we're going to follow Chrome and just say we're not going to honor any certificates that are signed. And certainly I would hope that CNNIC will do what they can to satisfy Google. Then Google will say, okay, fine, we're satisfied, and they'll be able to go back to issuing certificates.

MIKE: I also wouldn't be surprised if the Chinese government started a project again to - they're always trying to launch the Chinese OS, the China OS and stuff like that, and the Chinese browser. They may want to try to go it alone because they don't like being in this position where a foreign company can do something like this and have power over a Chinese authority like this. So I wouldn't be surprised if there's a newly revived effort, especially under the current Chinese leadership, to at least get Chinese users using something other than Chrome.

Steve: Yeah, I think that our topic for next week is going to be something know as the "Chinese Cannon." We're all familiar with the firewall. It turns out that this thing that was doing the big DDoS on GitHub and other sites has apparently some very strong Chinese backing. And from an analysis of it, it looks like there is an actual technology, a device technology that has been dubbed the "Chinese Cannon." And I want to go into the operation of that next week. So, I mean, we really are seeing an escalation of tension on the Internet.

MIKE: Yeah. And as I understand it, and again, this is from a news perspective rather than an examination of the actual technology behind it, but my understanding of the "Great Cannon" is that they're taking incoming traffic from outside of China that's intended for Baidu, and they're redirecting it at a target like GitHub, like GreatFire.org. And this apparently mimics to a certain extent NSA and GCHQ technology, which does something vaguely similar as part of a counterterrorism and surveillance effort. And so I can't wait to hear your treatment on that next week because it's really a fascinating - this is the first time that, according to some of the people quoted in some of the, you know, Wall Street Journal and so on, that it's been sort of legitimized, this idea that you are rerouting traffic to do essentially DDoS attacks in order to censor the Internet outside of your authority. So that's the part of it that's so fascinating is the Chinese government is censoring American websites.

Steve: Right, yeah, exactly. They're able to generate essentially a reflection attack against properties, as you said, outside of their control. Wow.

So a large botnet, which was known to have enslaved more than three quarters of a million machines, 770,000 computers spread across 190 countries, known as the Simda, S-I-M-D-A, Windows botnet was brought down through a collective effort a couple days ago. This thing was very aggressive. It was grabbing about 128,000 new machines per month. It was being monitored by a number of people.

And, for example, I think Kaspersky has a list of the IPs which were known to be infected because they ended up tracking down the 14 distributed command-and-control servers located in The Netherlands, in the U.S., in Luxembourg, Poland, and Russia, and did a coordinated takedown of those 14 servers. Well, if they're able to monitor those servers, they're able to get the IP addresses of all the infected machines which are phoning home to those servers for botnet instructions. And so that's how they know how many machines are infected and the rate of infection. The U.S. had the highest rate at 22%, almost a quarter of the total infections, followed by the U.K. and Turkey at 5% each, and Canada and Russia at 4% each.

MIKE: We're No. 1.

Steve: Yeah, huh, exactly. And the infections were exploiting known vulnerabilities. So these are machines not being maintained, so probably not a large percentage of this podcast's or even the Twit Network's listeners, because these people are going to tend to be more tech savvy. But so they were exploiting known vulnerabilities in Java, in Adobe Flash, in Microsoft Silverlight, in order to get themselves in. And what was installed was a highly stealth backdoor trojan which was morphing itself every few hours in order to be essentially polymorphic, as the term is, in order to make itself continually undetectable from existing AV tools.

And what's really sneaky is, you know, we've talked about, just in the last couple weeks, we were talking about how websites that - in fact, it may have been in reference to the Chinese DDoS attack because remember that they were using, in that case, they were using the Baidu Analytics. The Baidu system has an analytics system much like Google

Analytics, where all sorts of websites all over the place have a little bit of JavaScript which is added to their page, which causes the user's browser to query and pick up that JavaScript, that then feeds statistics back to headquarters in order to generate stats, in the same way that Google Analytics does.

Well, what this trojan was doing was it was modifying the Windows hosts file for connect.facebook.net and google-analytics.com. So where we were talking last week was this notion of all these little Facebook icons, the Facebook "like" icons that are everywhere now. All of those essentially phone home. They pull from connect.facebook.net. That's where that little icon comes from. So what that means is, if you get something in your machine which modifies the hosts file - as we've discussed before, the hosts file is the first place that Windows goes for DNS lookups. Before it asks any external DNS server, it looks in this so-called hosts file that Windows still has. And this hails from the original UNIX hosts file, which is where all of this Internet technology came from. And any domain which is present in the hosts file is looked up in that file.

So the idea was that this backdoor trojan for the Simda botnet would put the IP address of one of these 14 command-and-control servers in the hosts file intercepting lookups, the user's machine's lookups, for connect.facebook.net and google-analytics.com, meaning that any website that the user went to that had a Facebook "like" icon or that had google-analytics running on the page in order to generate statistics and Google Analytics for that site, would actually end up connecting to the botnet and open a vector of control and potential infection. So this was big and nasty. And what it took was a joint task force of both private security companies and law enforcement to decide, okay, we've found them all. Let's take them all down. And that botnet is now out of action. So, whew.

MIKE: Nice.

Steve: Yeah.

MIKE: Unbelievable. That is incredible.

Steve: Now, there's a so-called "Rootpipe" is the name that's been given to this, a privilege escalation bug, which was responsibly reported in early October of last year to Apple by security researcher Emil Kvarnhammar. I think I got that name about right. And he named it "Rootpipe." He discovered the problem because he was going to be going to a security conference, and basically he needed something to talk about. So he just thought, okay, I'll just poke at OS X for a while. And he did some binary code analysis and ended up finding this. And just because he sort of wanted to. I mean, he didn't stumble on it. He said, "I need to find something to talk about." So he plowed in and found this thing, told Apple about it. Apple said, "Thank you very much."

Now, they're all very low key. They don't tend to disclose the way, for example, Microsoft does. Apple just says, oh, yes, we fixed a few things; and, you know, don't worry about it. So he tweeted on October 6th, a few days after talking to Apple. He said, "Details on the #rootpipe exploit will be presented, but not now. Let's just give Apple some time to roll out a patch to affected users." So he did the responsible thing. So it got patched with this most recent update that we had a few days ago.

But, somewhat controversially, Apple has said they're not going to patch earlier versions of Mac OS X. They said that it required a lot of changing around of things and that it's just not worth their time. So this has angered some people who for whatever reason aren't staying current with the latest version and are on earlier ones. I have seen some documentation of an alternative. And I was first thinking that I would provide links and some coverage, but it is a mess. I mean, it's, like, not the sort of thing you want to go off and do yourself.

So the problem is the classic, or, well, the problem is that what Emil found was a way for software you run - now, this is not some guy in Romania is able to take over your machine remotely. This is a local privilege escalation bug, meaning that it needs to be software running in Mac OS X which, if it's aware of this, that is to say, today that means if you have an older version of Mac OS X, it's able to give itself admin privileges if the account has them. So this is the classic "Do not run with admin privileges." Of course, for convenience sake, most people do.

So what Emil's own explanation in his disclosure said, do not run the system under an admin privileged account. Create a second admin account, maybe named admin, log in as the admin user, and then remove administrative privileges from the normal account that you normally use, then log out and go back into your normal account. If you are a normal user, then there's nowhere for this to escalate itself to. It can only escalate to the maximum privileges you have. And if you've neutered those because you're not running as an admin, then there's nothing for it to do.

Again, this is fixed in 10.10, but not in earlier versions. So, and I'm sure anybody who needs to who's a heavy Mac OS X user can find additional detailed instructions on removing admin privileges, if they're deliberating using an older version of OS X and choose to keep using it rather than staying current.

MIKE: Now, Steve, is this a good idea anyway, just as a matter of course, to not...

Steve: Yeah.

MIKE: ...run under the admin? Because this kind of thing, you know, if you're under an admin account, of course, you're essentially open to do admin-privileged things. Not just you, but malware, whatever. So you've already - you've been recommending that people do this for a while.

Steve: Yes, it's standard practice. And it's worth noting malware had been found that knew about this. So this is a perfect example of, just because a security researcher has found it and responsibly disclosed it, doesn't mean that it wasn't found by bad guys or maybe gray hat guys. I mean, you know, we presume that the NSA has all kinds of ways of getting into our systems, if they want to, because they're using things that haven't yet been discovered by anybody else and responsibly disclosed so that the manufacturer can resolve them. And in fact, we've heard that Microsoft tells the NSA about things they haven't patched yet, which is a little disturbing. I'm not so sure why that happens.

But, yes, standard operating procedure is that, I mean, and it's uncomfortable because it will get in your way. But it's the only way to be safe, and that is, remove admin privileges from the account you normally use. You normally only need them when you're installing software. And so if, I mean, if your life is, like, being a member of the press, installing software all the time, then, okay, you don't have any choice. Or actually there you really want to set up a VM and just use rollback, you know, set up a static image with a virtual machine and install stuff all the time in there and then just erase it when you're done rather than uninstalling things that you've installed. It's just, you know, not safe to do.

But so my point is I can see situations where somebody just by virtue of the mode they're in has to be installing stuff all the time. It may be less practical to need admin privileges there. Although you can oftentimes just run as administrator during the setup and provide the administration password at that time in order to give that installation process admin privileges. So it can also be made feasible. But, yeah, this is another example of where you don't want to be in admin because it was possible for installation software behind your back to get those privileges and get up to some mischief with them.

MIKE: Good to know.

Steve: There's one ransomware, you know, we've talked about CryptoLocker and that ilk. There's one called CoinVault. And one of apparently several CoinVault servers were found, and the keys were extracted from it. Kaspersky got their hands on the keys and has a service now called "No Ransom." So it's noransom.kaspersky.com. They don't guarantee - first of all, this is only for CoinVault, so it won't help you with CryptoLocker. It's not like any - it's not a magic bullet for solving this problem for everyone.

But I wanted our listeners to know, if they or they know of someone who has been bitten specifically by the CoinVault ransomware, Kaspersky may have the decryption keys for you. They didn't get them all, but they got a chunk, enough that it offers some hope. So I wanted to make sure everyone knew: noransom.kaspersky.com. And you pulled the page up right there on the screen, so that's what it looks like when people go there.

There's been just another piece of news I wanted to share. Nothing really coherent yet in this ongoing struggle that law enforcement and private industry have over encryption. Everyone knows that several years ago I decided I was going to suspend my work on the product I was excited about. I was going to do a really high-performance, you know, GRC-style VPN that would do all kinds of tricks in order to solve the problems that VPNs often have of not being able to establish a connection because of VPN-blocking software, and a lot of other features. I called it CryptoLink. I got a trademark. I got the domain names.

But then I could feel, I mean, even before Snowden this was, it just felt to me like we were entering an era of tension where there were some early grumblings about whether the government was going to allow there to be encryption that they could not crack. And of course then we've had the Snowden revelations that revealed how much of this surveillance was going on. So there was a piece in Ars Technica where one of many proposals was discussed, and I thought our listeners would get a kick out of it, yeah, the split keying. Just two paragraphs from the article. I think it was - was it Dan Goodin who did the piece? I can't see it on the screen there.

MIKE: Yes, it is, Dan Goodin.

Steve: Yeah. He does great reporting.

MIKE: Yes.

Steve: He said: "Critics have also raised concerns that any type of portal that gives government officials access to encrypted contents has a strong likelihood of backfiring. They said criminals or spies from hostile countries may exploit the weaknesses to obtain classified or confidential data, possibly on a mass scale. The split-key approach floated by Rogers" - who is the admiral or general or something, he's the head of the NSA, and he's been outspoken on saying, you know, I think he said something like, "We don't want a backdoor, we want a front door with lots of big locks on it." Of course, he wants a means to have those keys.

MIKE: And that's Admiral Michael S. Rogers with the NSA. And he's said a lot of controversial things, and unsettling things, in my opinion.

Steve: Yes, yes. So Dan writes: "The split-key approach floated by Rogers, for instance, requires a complicated system to allocate the keys, deliver them to each involved party, recombine them when a legitimate court order is issued, and destroy" that key once used. And so our friend Matt Green at Johns Hopkins, we're referring to Matthew all the time. He of course, we were talking about, he was a subject of last week's podcast because of the TrueCrypt audit that we covered in detail. So Matthew said, quote: "Get any part of that wrong, and all your guarantees go out the window."

And so, continuing the last paragraph from Dan, he wrote: "The approach is only one of several options being studied by the White House. One alternative under consideration would have a judge direct a company to set up a mirror account so that law enforcement officials conducting a criminal investigation could read text messages shortly after they are sent. To obtain encrypted photos, the judge could order the company to back up the suspect's data to a server while the phone is turned on and its contents are unencrypted." And then the article ends: "White House aides hope to report to the President this month."

So, lord. I mean, I don't know. This is - all I can say is we're living in interesting times because, I mean, this is binary. Either the system will not have a means for third-party, under any circumstances, court order or not, it simply will not or it will. It's binary. It's not gray. So it either will or it won't. And the question is, are we going to see in the United States legislation mandating the binary decision that, yes, anything encrypted needs to have a means of court-ordered decryption, or are we not? It's one or the other.

MIKE: I have an overarching theory about the evolution of technology and its political consequences. And typically the general theory is that, every time there's any sort of advance in technology, people in positions of political power will take advantage of that advance to change the balance of power between the public and the government, essentially.

Steve: Yeah.

MIKE: And so this is a case, if you wanted to compare this to real life, this would be a case where the NSA would require that every single person wear a microphone 24/7 so that, even if they're out in the middle of the street whispering to somebody, that that could be accessed by the NSA. And that's just my opinion. I think that there has to - I think that, like real life, our digital lives, we have to be able to whisper things to each other sometimes and have secrets and privacy, just as a fundamental human right.

Steve: Yes, well, in fact, I think the best example of this is the argument that people say, why do you need privacy if you have nothing to hide? And so the response is, okay, what about a camera in the bathroom? Who's going to feel comfortable with that? They have nothing to hide technically that's presumably going on in the bathroom. But you have a right to privacy. I mean, the argument is that the government shouldn't be able to see into our private lives where it's just private.

MIKE: And that argument also, by the way, has implicit as part of an assumption behind that, that law enforcement and spy agencies are completely trustworthy, they never break the law, they never violate the Constitution, they never do any of those things. But of course we know with absolute certainty that there is lawbreaking, that there is abuse. And so there really is no authority that is going to be perfectly trustworthy in perpetuity. We do have to protect ourselves against official abuse.

Steve: Well, and of course anybody who has seen John Oliver's interview of Snowden understands the problem. I mean, Snowden absolutely knew from first-person experience that there were a lot of naked photos being passed around the NSA because they got them, they had them, and they're just human. I mean, and so there is a problem when humans are being asked to police other humans.

MIKE: Yeah, absolutely. Absolutely true.

Steve: So it was either last week or the week before that we were talking about cipher suites. And Leo put in Bank of America into SSL Labs' site and, in real time, on the air, we were stunned to see the ordering, the preferred ordering that Bank of America chooses, the idea being that, when a browser connects to a server, it presents the server with the list of cipher suites, the encryption algorithms, the encryption hashes, the encryption strengths, whether it's RSA or Diffie-Hellman, the key exchange technology.

Basically all of those are bundled up in a set of standards which each represents a suite of encryption technology. The browser says, here's the list that I know of. Then the server, the way SSL and TLS handshaking works, the server looks at the browser's list, and then from its own ordered list, meaning from most desirable to least, it uses its list to pick the cipher suite that'll be used. And so all logic says that the server puts the strongest ciphers at the top and works its way down to the weakest ones because the logic in the handshake, the logic at the server side is it looks at the first one it knows about and looks for it anywhere in the browser's list. Is it there? No. Okay, now the next one, and looks for that anywhere in the browser's list, is that there, and so on.

Well, it turns out that Bank of America, which was the only banking site we looked at, has the worst of all, like RC4 cipher, which has been roundly criticized and deprecated, with a short key length and bad encryption and a bad hash, I mean, it's just everything about it is wrong as its No. 1. So if a browser offers that, just for the sake of compatibility, the browser doesn't ever want that to be chosen. But it would rather that to be chosen than nothing. So the browser might have that properly at the bottom of its list, not that the browser's list is ordered. But basically it's in the list so that it's a possibility. But it doesn't want it chosen. Yet Bank of America sees that among the choices and chooses it first.

So we were just stunned. I just wanted to acknowledge that many people checked their own banks, and apparently this is universal. I got a whole bunch of tweets after that saying, you know, people were just like, oh, my lord, my bank is worse, is as bad as BofA. So they were just tweeting their despair over discovering that their own banks were in as bad a shape. And in fact, a question that I saw in the mailbag, I didn't choose it, only because we've sort of covered it, and I just have, was somebody saying, how do I explain this to my bank? I mean, you know, I'd like them to fix this. But who? How? You know?

MIKE: Walk up to the teller. Excuse me, I'd like to mention something.

Steve: We have a problem, you have a problem with your cipher suite ordering on your website. So, yeah. I mean, I don't know how. Just maybe - I don't know. Maybe there's a support email on the website, you know. Maybe, if they get enough email from people saying that, somebody - it'll, like, percolate up. I mean, there is no fathomable reason for that list to be as it is, except nobody ever curated it. I mean, in fact, it's hard to understand how the server could have even been shipped that way. But there's no reason not to put that at the bottom because, if no better cipher matches first, then it'll still be chosen. But don't pick it as your first priority. That's just nuts.

MIKE: There's so much bad decision-making and bad design in websites for this exact reason, assuming that this is the reason, that is, that something was assigned to somebody who is not looking at the big picture. They're just trying to go through their list, their checklist of things to do.

Steve: Get it done, yup.

MIKE: And nobody pays attention, nobody follows up, nobody audits it, nobody, nothing, and it's just there forever. And it's just a terrible state of affairs.

Steve: Yeah. And unfortunately, as we know, the nature of security is a weakest-link-in-the-chain model.

MIKE: Yeah.

Steve: That is, in order to have security, every single piece has to be secure. And, I mean, I don't know how this ever changes. You know, I've used the example over the years of firewalls, how the original design of firewalls was that it was all open, and then when something bad would happen, you'd close the port that that had been done on. And so over time you'd close more ports as more mischief was being done. Well, we've completely reversed the model now so that the firewall, all ports are closed, and you only open the ones that you know you need because it finally occurred to people, okay, we've been doing this wrong. And the problem is it's not clear how we can get to the same sort of model with security.

As a consequence, if a website is not designed with every aspect of its security in mind, somebody will find, even with an automated tool that just goes scanning around websites, looking for problems, somebody will find an opening and compromise your web server and then use that to get into your network. And then you've got another Sony Entertainment catastrophe.

MIKE: Yeah, very true. And usually the best way in is through users. We covered on Tech News Today this morning two new reports from Verizon and Symantec, separate reports, both of which said that one common feature of these big hacks that we've been reporting on all year, this year and last year, is user error, usually phishing attacks. Phishing attacks is the door to get in.

Steve: Yup, yup.

MIKE: And it's difficult because you can either shut down attachments altogether and sort of hobble the usefulness of email or whatever, or you can train people. But no matter how much training you do, the user is always going to be the weakest link. And so that is, I think, a huge challenge. And it's interesting to see these two reports say the same thing.

Steve: Yeah. And in fact we've discussed, it was a Q&A two weeks ago, and we talked about it again last week, it was a fun question that someone said he was in charge of security for something like North and South Dakota something or other, I don't remember now what, or the north and south of something. Anyway, he was in charge of security. He had managed to convince his users never to click on links in email. And he said, "But when I send them links, I get all this flak back now, which I guess is good because it means they've learned what I've been trying to tell them, except it's from me."

And so of course we explained that, yes, unfortunately, that many phishing attacks appear to come from your mother. And so they don't because Mom's AOL account or Yahoo! account got breached, and something evil sent email out to everyone your mom knows saying, oh, this is the funniest cat video you'll ever see, you know, click here. And because it comes from Mom, it's like, well, of course that's innocent.

MIKE: But, you know, I think one of the most shocking things from our report today, we talked to Joseph Menn, we interviewed Joseph Menn, who's a reporter for Reuters, who covered these two studies. And he said that, according to one of them, I think it was the Symantec study, you need to send 10 emails, only 10 emails, to have a 90% chance of success with malware.

Steve: Wow. Yeah.

MIKE: And of course it's free to send 10 or a hundred or a thousand.

Steve: Wow.

MIKE: Incredible.

Steve: Yeah. And, well, in fact we know that the massive RSA breach that we covered years back where RSA had their internal network, it was a persistent threat that had been in their network for a long time, and a third-party security firm was brought in, they did all the forensic studies, and they found an assistant had clicked on a link in a phishing email, and that's the way this APT, this Advanced Persistent Threat, got into RSA and to obviously devastating effect.

MIKE: Yeah.

Steve: And as far as we know, same deal with Sony. I mean, you just can't, unfortunately, you can't click links in email. And it's interesting, too, because there was a lot of feedback in this week's mail bag about people's different suggestions and things they tell people. And what I said last week, following up on this, was the only thing I could see that this guy could do is, if he was in charge of security within some organization, he could maintain a web page and tell people to manually type - make the URL simple so it's news dot security dot whatever it is dotcom, or securitynews dot whatever dotcom, just go there, and there you will find the link for the asset that I want you to click on, the idea being - and this is one of our other standard truisms is rules of conduct. Never, never do something that is offered.

That is, for example, if you go to a website, and the website tells you your Flash is out of date, do not install it. You didn't go seeking it. It is seeking you. And so don't click on something in mail to go somewhere. Instead, you go there yourself. So the idea being, you know, and I think this may originate with Brian Krebs. If not, we saw it there first, and that was don't install anything you didn't go looking for. If it's offered, just decline.

MIKE: But again, that's a training issue. And I don't know, have you covered the Dyre Wolf on this show?

Steve: No.

MIKE: Okay. The Dyre Wolf is something that IBM discovered recently, and it's a sophisticated fraud that combines phishing, malware, and phone calls. And they've extracted a million dollars from U.S. companies, which is not a lot of money. They kind of nipped it in the bud. They're calling it the Dyre Wolf, but it comes from Eastern European cybercriminals. And what these criminals were doing was they had a - it started with a phishing attack, malware attack. So they get this stuff running on people's systems, employees within banks. Or, no, I'm sorry, I've got that wrong, a bank customer.

So they get it running on the side of bank customers. And it would sit there and lurk in the background and wait for the customer to go to the bank website. And it would see the bank website, and here's the genius part of it. It would identify the bank, grab the graphics from the website, and put up a fake website that looked like it was part of the bank that says, "We've got a problem with your account." Okay, all this stuff is the kind of thing you'd expect; right? If there's a problem with your account, we need to talk to you about it. But don't do anything on the web. Give us a call. Here's our number.

Steve: Whoa.

MIKE: And they had criminals standing by, and the malware would actually send the bank information to the operator so that, when they answered, they said, "Hey, Bank of America, how can we help you? Oh, yes, yes, I see that your account is..." thing. But, I mean, what do you do with something like that?

Steve: Wow. Wow.

MIKE: That is really sophisticated and very deep pockets kind of an attack where they're actually hiring a call center staff.

Steve: Yup. And it also incorporates the whole social engineering side.

MIKE: Yes.

Steve: I mean, it's one thing to just sort of click on a link that says you need to update Flash. But as you say, this ups the ante because now you have what you're not normally going to get, which is, oh, I dialed a phone number, and I'm talking to a person.

MIKE: And the fact that they told you don't do any of this over the Internet engenders trust.

Steve: Yes, and much more credibility. Wow.

MIKE: Yeah.

Steve: So two miscellaneous bits. Tonight, thanks to all of our listeners, with any luck, although we know luck was not involved, Security Now! may be the winner in the Technology category of the 10th Annual Podcast Awards. For anyone who's interested, that ceremony begins streaming live over PodcastOne.com beginning at 6:00 p.m. Pacific time, so that's 9:00 Eastern. And we'll see who the winners are of the 10th Annual Podcast Awards. Leo just sort of rolled his eyes when I said I wanted to do this. I don't know, there's some history there.

MIKE: Well, Leo doesn't like to do any sort of explicit promotion like that, self-promotion like that. I tend to be highly self-promotional. But Leo just - he just wants things to happen as they're going to happen, for the most part, and doesn't like this kind of explicit thing. But, hey, if you win, that would be awesome and well-deserved.

Steve: Well, I've got two trophies, I have two trophies behind me from prior years. And I haven't, we haven't done it for many, many years. And I just thought, okay, this is the last time, just sort of as a show of strength. Let's see if we can mobilize our listening forces and make it happen. So I wanted to let everybody know because they closed on the 26th, I think it was, and now they've been counting ever since. And tonight we'll find out. And I'll mention who won next week.

And then the second note is there is a TV series that I turned our listeners on to several years ago that has been a huge hit within this community. So I wanted to make sure, thanks to Randy Thomas, who tweeted from Rapid City, South Dakota, or maybe, actually, maybe it was in my email pile that I encountered, Season 3 of "Orphan Black" resumes this Saturday, April 18th. So I know we have a huge following of "Orphan Black" fans who follow the podcast, and so this coming Saturday. I checked, and my TiVo already knew. It was in the queue, ready to be recorded. So it will be interesting to have that adventure continue. So I wanted to make sure that our "Orphan Black" fans didn't miss the beginning.

MIKE: Very cool.

Steve: And I have another fan, in this case of SpinRite. We have a French language speaker. He said, "Hi, everybody. First, I want to apologize for my imperfect English." To which I always reply, well, it's way better than my, well, nonexistent French. So I never have any problem with somebody speaking my language less perfectly because I don't speak theirs at all. He says, "I'm French speaking. I want to share with you my SpinRite story. I purchased a license to SpinRite, mainly to support you. I'm a big fan of Security Now!. Anyway, I never got the chance to really test my version of SpinRite since I bought it. I mainly use the program for maintenance," and he says, "preventative scanning." And those obviously are of his various PCs.

He said, "A couple of weeks ago my Cisco Explorer 8300HD DVR," he says, "cable TV set-top box started to behave strangely." And you might want to put up the screen that I have on the next page, Mike, because this is where he's leading, that is, the next page of the notes. He said, "...started to behave strangely." He said, "A lot of recorded shows would pause and skip frames. In worst case, I got some shows that would not record at all. I suspected the hard drive to be the problem, but since these devices runs on custom "encrypted" file system, it was impossible for me to run any usual programs like Windows ScanDisk.

So I plugged the DVR's HDD into my main PC and ran SpinRite at Level 2 on it. When I checked the S.M.A.R.T. screen," you know, S-M-A-R-T, "I immediately saw the problem." And then in the email that he sent, he sent this screen, which is here in the show notes. He says, "The HDD," the hard disk drive, "was slowly dying. I was lucky enough, the faulty HDD was the external one, the expansion. SpinRite fixed it, and then I replaced it. And since I can enjoy my recorded TV shows and feel confident" - oh, he says, "And since then I can enjoy my recorded TV shows, and I feel confident about my scheduled recordings. Another example of how SpinRite can help in some unexpected situation. Throw any disk at it, and let it do its magic." I think his English is just fine. He says, "Thanks, Steve, and all who work with you on SpinRite. Steve Rodrigue."

So, Steve, thank you for sharing your story. And what the screen shows, for anyone who's interested, is the scary thing is that red dot under, on the very top line, showing ECC corrected. What's happening is that those three lines are three S.M.A.R.T., Self-Monitoring Analysis and Reporting Technology is what S.M.A.R.T. stands for, those are the three parameters which this particular drive publishes. Not all drives publish the same things. They're just based on make and manufacturer and what they feel like bringing out to the surface. But those are the parameters that they publish.

The problem is that they only mean something when the drive is under stress, that is, when it's being used. If the drive is just sitting there, you're not asking it to do any reading or writing or seeking or anything, then those parameters really don't mean much. So it turns out that the use of S.M.A.R.T. and SpinRite are synergistic because the idea is these should not fall even in the presence of asking the drive to do anything. After all, the drive should be able to read and write anything you ask it to. That's all SpinRite is doing, essentially, is really giving it a workout, while turning off a lot of the things that tend to cover up problems so that SpinRite can see them and so that it can show them to the drive.

And so in this instance the ECC Corrected parameter has dropped, meaning that while SpinRite was reading the drive, the drive itself was being surprised by how much error correction it was being required to do. SpinRite was able to get the drive to do its correction. If the drive saw that a sector was worse than it expected, it would have replaced it, thus fixing the drive. But ultimately a drive that's having this much trouble is screaming out, saying I need help. And in fact the second, the lower half of this shows the total error correction count, which is a little on the high side, 2,321,837 up to whatever point this was. And then SpinRite also shows the minimum per megabyte of data written and the maximum. And that's another indication of a problem. If there's a huge variation between the error rate, then that's a problem.

So here we see a minimum of 4300 and a maximum of 45, almost 46,000, meaning that there were some areas that were low, and others like rough patches. You don't want rough patches. You'd like the minimum/maximum to generally be pretty close to the average. In this case they weren't. So lots of indication here that, you know, it's not clear how much more life this drive has. And in this case, Steve Rodrigue was able to have SpinRite fix it, and then he took it out of service gracefully after watching all the programs that were on it.

MIKE: Unbelievable. Another SpinRite success story. Well, it looks like we are ready for the Q&A portion. I love this. I've been reviewing these questions, and I can't wait to hear what your answers are, frankly. This is listener-driven potpourri No. 210. And if you're ready, we can jump right into it.

Steve: You bet.

MIKE: All right. Well, the first one is a very, very basic question from Scott Nickles, who asked it via Twitter. And here's what he asked: Can you help me understand this? How does WiFi assist location info on cell phones? Of course, as we know, there are many ways for a smartphone to detect location, and one of them is WiFi. But Scott is asking, how the heck does this work?

Steve: Right. So I guess, what, all of the radios can potentially be used for location services, meaning that many smartphones now have GPS. So if there's GPS, the global satellite positioning system, then that is obviously a way to get located. And then, famously, and of course fodder for so many spy and techno movies, is the cell tower location, where we know that the way cell phones operate is that there are, scattered around, individual cell towers; that the whole concept of cellular technology is a multitude of short-range radio connections, rather than just like one monster antenna up on a huge hill somewhere that all of the cell phones talk to. And so, clearly, it's possible to locate someone to some degree, first off by knowing what cell tower they're using.

But it's often also the case that multiple towers will be able to receive the signal from a single cell phone. And so to some degree, by comparing the signal strength or the power level that the various towers are seeing, you can even get better than, like, what is the closest tower. But, for example, there's a radio sort of midway between two, which would be the case if they both were seeing about the same amount of radio energy from the cell phone. That's still subject to a lot of uncertainty.

What's interesting, though, is that, thanks to all of these crazy street view applications and mapping, one of the things that's happened is that WiFi routers everywhere have had their locations mapped. That is, of course, this famously got Google in trouble in Germany when they were doing their Google street view and mapping technology, they were sucking in all of the WiFi signals of the routers in their immediate area. They also had GPS so that they knew where they were. And so what Google was doing was deliberately mapping which WiFi routers, sometimes by SSID, or by MAC address, because MAC address is supposed to be universally unique globally, and that's easily seen. So Google built up a correspondence between location and router and router signal strength. So that essentially, every single router that is within a region, its MAC address and location is known.

And I'll never forget, early in the iPad days, I had an iPad with no cellular radio. And I was having breakfast somewhere and went to the maps application, and it just nailed where I was. There was that stickpin sticking up with a little sonar circle going bloop, bloop, bloop. And, I mean, it was exactly where I was. And this was in the early days of this, and I was dumbfounded that I was located, with no cellular connection at all, exactly. It's because I was within range of a standard WiFi hotspot, and the MAC address and probably SSID of that hotspot, its location was known exactly. And so through this system of associating locations, the exact locations of WiFi hotspots, my location was known. And all of this technology is now, you know, databases are available, and it's completely available to cell phones.

MIKE: And clearly Apple is willing, too, because every time, oftentimes I turn off WiFi because my phone is trying to get a WiFi connection, and I just want it to go - I know that the WiFi is off the table, so I want it to just go to mobile broadband. And it says, oh, wait a minute, are you sure you want to do this? Your location needs the WiFi. And I always think, wow, that's amazing that it's important enough to Apple to actually throw up an error message like that, just when I turn off WiFi.

Steve: Yeah. Even, I mean, and these are not hotspots you're connected to. That's the other point is in the same way that, when you are using a laptop, and you go somewhere, and you say show me a list of WiFi. Well, there's often 25 things. I mean, there's all kinds of crazy WiFi within range all the time. And so, exactly as you say, Mike, it is so important to location because now there's a map of known router identities and location. And I'm sure they're checked to make sure that they haven't gone bad, that somebody in Arkansas hasn't moved to California, and now their router is suddenly in the wrong state, because you would have to track that. But the point is it's not even necessary to connect to these. All of that information is just available in the air to any WiFi receiver which is just sniffing whatever happens to be available.

MIKE: And this is one of the reasons why I was critical of the Germans for attacking Google over this point.

Steve: Yeah.

MIKE: Because they made it seem like Google was reaching into people's homes and into their routers and routing around and looking around. No, you're broadcasting that information out into the street. And it's like putting a TV set in the window with your password written on it, this electromagnetic radiation conveying public information out into the world. And if somebody wants to come along and take a picture of that, it's identical to what Google was doing with the Google street view cars. But I was probably in the minority. People didn't like it. And it was all a big mistake or task or trial or something that was unauthorized anyway, apparently. So what are you going to do? Very interesting.

Steve: Well, you may have been in the minority, but I was with you because it was clear from - I remember looking at the software that was used, and it was clear that the engineer that put it together got that as a side effect. They weren't even, I mean, they didn't care about what the content was. And if it had been encrypted, fine. The fact that it wasn't, wasn't their fault. And the software logged this by default. So it wasn't something that was turned on in order to be sneaky. It was just, I mean, they weren't even using the information. So I'm with you. This was a mistake, and this is something I think Google got into a lot of trouble with that was really unfair to Google.

MIKE: Yeah, absolutely. Okay. So our next question, No. 2, from Gail Standig-Portman. She had some info about the Configuration Mania Firefox extension. And she writes: Steve, in your last podcast you mentioned the Firefox add-on Configuration Mania. I tried it; but it reset most, if not all, configuration changes I had already made. I've just gone through them again, trying to reset them back to where they were. You have a terrific podcast, but please warn your other listeners of this problem. If someone wants to use the add-on anyway, it would be very helpful for him or her to make note of any previous configuration changes so it'll be easier to make those changes all over again from Configuration Mania. Best wishes, Gail, a long-time listener.

Steve: So just as a public service, I wanted to absolutely make sure everyone knew. Another listener wrote in, or maybe tweeted, about Configuration Mania because I was talking about this overwhelming list of options you get if you do about:config in Firefox. I mean, it's astounding how many little tweaks there are now, so much so that you can't scroll through it, that you need the search bar there, and you type, for example, "OCSP" if you're looking for something having to do with certificate verification on the fly. And it'll then reduce that to three items, and you can see what it is you want.

Anyway, so it turns out Configuration Mania, come to find out, thanks to Gail, when you install it, apparently it doesn't first read the existing set of configuration, if that's even possible. Maybe it's not available. But they certainly should warn you that any configurations you have made are going to be reset to default because this thing can only write to the configuration file, is unable to read it, presumably. I mean, it's hard to understand how it would operate that way, but I absolutely wanted to make sure that no listeners got, if they would spend a lot of time, as it sounds like Gail has in the past, fine-tuning her Firefox config, and boy, there is plenty in there to tweak, I wanted to make sure that no one else got zapped.

MIKE: Yeah, excellent. Okay. The third note here comes from John Reagan in Nicholasville, Kentucky, who had a correction on HP's Class 'A' network. And John writes: Hi, Steve. In the last Security Now!, you said HP's Class A was 14. In fact, it's 15. And since HP bought Compaq, which had previously bought DEC, it owns 16, as well. That was DEC's Class A address. I wonder if HP knows how much money one of those class A networks is worth? John Reagan.

Steve: So I actually didn't intend to be correct last week. I knew that it was a low number like that, and I just sort of said 14. So I really apologize because I should be more careful. John is one of many people who corrected me. So for the record, because I do care about being correct, HP is 15.x.x.x. And I also knew that they had two consecutive ones. John provides the information that they got 16 from DEC because DEC used to be 16.x.x.x. And so I guess that's one good thing that Carly Fiorina did for HP is to double their Internet IPv4 space because, boy, those things are valuable.

We talked last week, Mike, about how there was a company, Jump.ro, I think that's - or "ru"? No, "ro," it's Romanian - that has a whole bunch of IPv4 space that it's not using. So they're selling them off and making millions of dollars because they're now becoming so scarce, and people are - and certainly it'll be an interesting shaped curve because the price is going to keep going up and up and up and up as IPv4 space runs out. And the pressure not to have to actually ever, please don't make us move to IPv6, grows. But at some point it's going to happen. And then suddenly IPv4 won't be worth anything. So, you know, there's a formula there. It's like, okay, I want to hold onto my unused IPv4 right up until it hits just before the decision made to finally give up on not holding onto IPv4 and capitulate and go to IPv6, where there are so many ridiculous IPs that everybody gets to have a Class A network of their own to play with, and good luck to you.

MIKE: It's like, so, yeah, absolutely. It's like ber surge pricing. If there's only one car left in town, you're going to pay for it one way or the other. Okay. So No. 4 we have from Matt DeWater in Auburn, Washington, who wanted to make sure about SpinRite and PCIe SSDs. He writes: Hi, Steve. I'm looking at doing my next computer upgrade and have been considering a PCIe SSD. Will I be able to run SpinRite on a non-SATA drive? Thank you for everything you do. Matt.

Steve: And I just wanted - I've seen a couple questions like that, so I wanted to assure everyone, yes. First of all, you'll be able to run it under SpinRite 6 today because the BIOS will understand how to read that, and SpinRite knows how to talk to the BIOS. What I'm doing in 6.1, which I promise to be back to as soon as SQRL is put to bed, and we're close on that, in fact just yesterday we eliminated a couple features that we decided we didn't really need, and so it got us that much closer. What I'm doing with 6.1 is essentially finally ridding myself of the BIOS. What that means is that I'll be talking to the hardware directly. And the good news is there's a standard called AHCI, Advanced Host - I'm blanking on the name. I know it so well, I can't believe it. It'll come to me. Anyway...

MIKE: Controller Interface.

Steve: Controller Interface. Advanced Host Controller Interface. Thank you. SpinRite, actually, where we left SpinRite was I was talking AHCI, allocating a huge amount, 32MB of RAM in real mode, which is a trick because real mode is supposed to be limited to that 640K barrier, 32MB, and transferring - we benchmarked it, transferring at half a terabyte per hour, meaning that it'll be able to run, do like a 2TB drive overnight, in eight hours, which will be a really nice performance boost. It can actually do that in some cases with today's BIOSes, if they've been written well. But that's the problem is BIOSes, many have just been neglected, and so SpinRite ends up having a problem talking through these older BIOSes. But so absolutely, you can purchase PCIe SSDs with confidence, and SpinRite will be able to keep them in shape for you and recover problems when they occur.

MIKE: Nice. Okay, No. 5, Mark Grennan in Oklahoma City wonders why self-signed certificates are not "TNO." He writes: Hi, Steve. You've been talking about SSL certificates a lot, and you've said a couple of times on the podcast that certificates signed by certificate authorities (CA) are more secure than self-signed certificates. I'm paraphrasing. But does this not break your TNO policy? I feel that generating my own certificate for my private email server is more secure. And wouldn't this be true for any private service, like company employees connecting to the company "X" service? I understand that, for the general public, adding a "public" CA cert for every service they want to connect to would be too much. And that is the question.

Steve: So, okay. TNO, of course, is the acronym I coined years ago, Trust No One, which we apply mostly to cloud stuff where, like with Dropbox, if Dropbox is a service for storing your files, and they say, oh, you're able to get your files when you're on the road, just go to the website and log in, and we'll give them to you. Well, that requires them to be able to do the decryption themselves in order to give it to you without having a client at your end. So there are ways around that, but that's not Trust No One. So Mark is applying the Trust No One test to certificates that are self-signed versus certificates where a third party whom we do trust has signed the certificate and is attesting to its authenticity.

In the example Mark gives, I would argue, yes. If you're a private company, and you want to create a secure connection to an internal server, and everyone you need to have trust that certificate can do so, then there's no reason not to do it. For example, the reason we use a third-party signer is that everybody in the world, for example, used to trust CNNIC. Well, now only half of the people in the world trust CNNIC because of their misbehavior. But the point is that the fact that there's global trust of the CA means that their signature extends that trust to the certificates they sign. So it is necessary to trust the CA.

And so the perfect example is CNNIC, whom many people no longer trust. If you sign the certificate yourself, then the problem is the rest of the world won't trust it. They'll look at it and go, eh, this is a self-signed certificate. Now, many people have gone to websites, and you're presented with, like, the notice, hey, this website has signed its own certificate. So you could trust it if you want to, but we're not implicitly trusting it because the certificate was not signed by somebody who verified the site's identity.

And that's the point is somebody could have gotten that site's self-signed certificate and redirected you with DNS to themselves somewhere else and be impersonating them. I mean, the problem is that, with a self-signed cert, you're trusting the person who signed it. So outside of a local, like an Intranet, a corporate framework, no one's going to trust that. Inside, with a private service, sure, you could say to your employees, hey, when you check email, and your email client comes up and says, hey, should we trust the certificate, just say yes.

And so you get a secure connection inside the corporate environment. You're not having to pay a certificate authority. Nobody outside needs to trust this service, and they wouldn't because the certificate is signed by itself, essentially, by you, not by someone that everybody trusts. So it's sort of, you know, a choice you can make about who you want to trust the certificate. And if it's local, then, yeah, you can get by with a certificate that you sign yourself. And you don't have to pay anything, and you can set its expiration date way into the future, also.

MIKE: You need to coin a new set of initials called N-O-T-U, or NOTU, which means No One Will Trust You. I guess that would be a Y. Okay, we'll work on that. All right. No. 6, Tom Etienne in Killeen, Texas, wonders about secure email providers in 2015. He writes: Dear Steve, it must have been a year ago I stumbled across Security Now! following announcement of the Heartbleed bug. It was fantastic to discover a way to hear a deep examination on these topics given the limited amount of time I have for reading and research on my own.

I've recently begun to investigate alternate email provider options, given the exposure of Gmail and other similar large providers to both internal and external tracking. While the features and convenience of Gmail are substantial, privacy remains questionable. I can adopt just about any new tool or service that is available, but I am by no means a security science guru with the ability to audit their stated claims.

Some new email providers have arrived this year, such as ProtonMail.ch and Tutanota.de, to join alongside outfits like Hushmail. I'm participating in the beta invites for these new sites, but wonder whether or not it is premature to shift the majority of my email communications through one or more of these services. Review websites offer raves and rants about every service in existence, but I'm curious if you have any deeper insights or recommendations to offer about secure email providers.

Steve: So Tom's question is a proxy for a question that is asked so often. I guess the problem is that people would like secure email. And so the desire is to have secure email. The problem is that email's underlying protocol never incorporated security. I mean, there are variations. One's called S/MIME, which is the formal secure technology, like the standards-oriented technology for adding encryption to email. And of course, famously, PGP and GPG and so forth. But those have just never quite gotten off the ground because the email protocol itself didn't require them. So they don't get used by everyone. And if they're not used by everyone, then you don't get end-to-end encryption.

So the problem with any of these so-called "secure email providers" - and, for example, after reading Tom's question, I thought, okay, let's just check in on ProtonMail.ch, just, I mean, because I hear about them all the time. People are asking me about that in particular, that one in particular. And so I go. And it's a very pretty-looking website. Lots of assurances about security. No technical details at all. So the first mistake they make for me is I can't learn what they're doing. And our listeners of the podcast know that, if they're not willing to show us exactly what they're doing, then we're stuck.

A perfect example is a couple weeks ago I talked - or a perfect counter example is I talked about miniLock.io, which is that fabulous little public key encryption add-on for Chrome and the Chromium browser and Chrome browser. And, I mean, it's fabulous. And there on the page was a complete disclosure of the way it works. And I was able to read that, know that he had done it exactly right, and represent that to our listeners. And it's, like, our listeners are crazy about it because they understand from me how it works, and it was done right.

Now, here's ProtonMail, for example, that waves their hands around and has a very glossy-looking, pretty website. But they're not telling me how it works. And the problem is they're also not making it clear that, unless every party in the email is using ProtonMail, then you don't have security. That is, they play up the connection that you have to their server and how that's secure, and they don't have the ability to decrypt your email, all that good stuff. Except the problem is, if your email is going to someone else who is not a ProtonMail user, and you haven't explicitly enabled a feature that sends that party a link which brings up a web page to allow them to decrypt the mail that you're sending them, then it's decrypted and sent as regular insecure mail.

The point is, this is all incredibly awkward because email is not fundamentally secure. It just isn't. And everybody wants secure email. It's a nice-sounding phrase, "secure email." But email isn't secure. The protocol never had it. And so I guess, you know, so I look at something like miniLock.io and say, hey, for those specific instances when I need to encrypt something, this is the way to do it. I can send somebody my public key. They send me theirs. I encrypt something for them. They can decrypt it because it came from me and was encrypted for them. It's solved. The problem is solved.

So if you want everything you send to be secure, now you've got a problem because the only way that happens is if you are sending email to somebody else who is a ProtonMail subscriber, that is, or who is using the same service, and if we believe their assertions. They're not telling us what their technology is. There's no description that I could find of how this works. So we have to simply believe their representations.

And the problem is, even if their heart is in the right place, they may have done it wrong. We covered a few weeks ago that service that was offering web-based encryption, but was not doing any authentication. Clearly, they were trying to do the right thing, but they didn't do the crypto right. Because they told us what they were doing, I was able to say, well, nice first try. They didn't do the crypto right. And then they fixed it once, and then they still didn't do it right. The point is, because of the full disclosure, we could tell and say, eh, it's not ready yet.

ProtonMail, who knows? Maybe it's secure. It's certainly only secure if everybody you're sending to is also a user, or if you use that option where it simply sends them a link, and then they have to go and retrieve the email, essentially, back on the ProtonMail site. But I'd feel more comfortable if a security-based email provider offered a technology whitepaper saying this is the way we implemented our security. Then someone like me, or anybody else who understands crypto, could look at it and go, they did that right. They probably did. Maybe. But it's just as likely that they made some mistakes. Which if they would show us, we could help them fix.

MIKE: Great point. Great point. Okay, No. 7, Rafael Beraldo in Brazil has a routable IP from his college. And he writes: Hi, Steve and Leo. Last episode Steve mentioned that Jump.ro, a Romanian company, is selling IPs like crazy because they're worth a lot of money right now. Older universities also often have a lot of IPs they're sitting on, and mine is no exception. But this week I noticed something I'd never realized before. Whenever I connect to my college's wireless network, I'm given an IP on the Internet, and not anything like 192.168.042.042.

While I understand there are advantages - I can run my own services on low ports, for instance - I'm still concerned about the implications of having a public IP address. I run Arch Linux and make sure there aren't any services such as SSH running when I'm on a public network. But on the other hand, my Android smartphone, which is connected to the same network, has a public IP that's easily scannable. That's crazy. I feel kind of scared that this is the case, even though the network's gateway probably has a firewall. So how should I feel about this? I mean, it is easy to find the address space given to a college and then scan all devices for vulnerabilities; isn't it? Thanks for reading. Rafael.

Steve: Yeah, I loved this because it was my own, in fact, we coined a term, a "Gibsonian response," where you just sort of shudder. And it's interesting that I have that response when I imagine that the network I'm on is publicly routable. I guess it's because I am so used to being behind routers, that is, you know, private routers, that the router has a public IP. I have a 192.168.something.something, or in my own case I use 10-dot as my network. And even at the datacenter where I have our co-located servers, all of those systems are 10-dot, and there is explicit IP mapping going on in order to make individual services available that I want to make publicly present.

And so the idea that a smartphone, especially, as Rafael says, his Android phone, the idea that it has a routable address means exactly as he knows and wrote, that a scan of the Internet is going to scan the device he's holding in his hand. And if there's anything flaky in it - and my lord, when is there not something flaky in these devices? We're constantly talking about, for example, how backdoors are being found in routers. Well, it's only recently that Windows finally had its firewall turned on - well, recently, XP SP2 - that the firewall was turned on by default. But initially it was there, but not even on. And we know that, if you take one of these machines and stick it on the Internet, it'll be taken over immediately. There's still Nimda and Code Red out there scanning, looking for machines.

So I love, and I just take for granted now, that there is a router which is inherently a hardware firewall. Nothing that is unsolicited, no incoming unsolicited traffic goes anywhere. Only conversations that are initiated from inside to the outside create a dynamic map through the router such that returning traffic has a place to go. It goes back to the initiating computer on the inside network. I'll tell you, I mean, were I in a situation like that, I would get one of these little mini routers. I know that D-Link has one. They're called travel routers. Or just, you know, you could use any router, actually. But for some reason I just think of a travel router. Let that get the public IP address, and then you use its WiFi and/or wired network that is going to be protected, and then you know that you're not subject to attack.

Boy, I mean, and even if the college has a firewall, what about all the mischief that other people within that firewall boundary can get up to? That is, you're looking at an Intranet. There may be a firewall on the border. But presumably all the machines within that campus environment can see each other and are able to scan each other. That would just - I'd absolutely be behind a router because they're like doorstops now. You can just find them lying around.

MIKE: And of course, if you get a Gibsonian response to anything, you should be looking into things, as well, out there in Security Now! land. Okay, No. 8. Jim Fletcher in San Francisco was left a bit lost about TrueCrypt and keyfiles. He writes: Dear Steve, I was delighted with the TrueCrypt audit results and your great explanations on the show. This is great news and just as you predicted. Does using keyfiles in TrueCrypt make the encryption weaker? I love the concept of keyfiles. I make up a unique keyfile, for instance, a picture, which will of course be a unique file, then keep it on a thumb drive. It's two-factor authentication, and the thumb drive can even be kept in a safe. Brilliant idea. But since it's such a poor implementation, does it in some way weaken the encryption of the TrueCrypt container or partition? Great show, great product - SpinRite - and best hosts in the world.

Steve: So, great question, Jim. There were four problems that the crypto auditors found in TrueCrypt. And I think it was the weakest one, the least concern, or the second to least. There was that the headers of the containers did not have strong authentication, they just used a CRC technology. That was one of the two weakest. The other was this issue of keyfiles, which was they were unimpressed by the way the keying material was merged in with the user's password to create the key that would unlock the header.

So, and the short answer to Jim's question is no, it in no way weakens anything. This was them being crypto purists, which I completely understand. What they explained was, because the keyfiles were being combined in a way that was not cryptographically sound, it would be theoretically possible for an attacker who was present at the time the partition was being set up, at the time the keying material was brought together, when the user said, "I want to use this photo along with this password to encrypt this container," at that moment bad guys could look at the contents of the picture and come up with a nullification, come up with a modification that would nullify the contents of the additional keying material that the picture represented. You could not practically do that if the picture was hashed because, if you're hashing a bunch of things together, the point is it is not mathematically feasible to come up with additional material that creates a given hash result.

And but these guys, the TrueCrypt authors, didn't use a hash, for reasons no one could understand. They used a simple cyclic redundancy check, a CRC. And it is absolutely possible to mathematically compute a way to reverse, like to create any output from a CRC that you want to. It's trivial. But the point is, even then you have to be present at the moment this is all being done. I mean, it's not even clear, they didn't articulate an attack. They just said we don't think this is very good. But they didn't even say how you could use the fact that it wasn't very good because you actually can't. But it's true that it's not very good.

And so this is really the only use for this is people who are forking the TrueCrypt source. They could look at the audit results, and I'm sure they have, and go, oh, yeah, we'll use a hash. Problem solved. So, Jim, this is nothing to worry about. You absolutely have strong multifactor authentication because no one was present, no bad guy happened to be watching when you set this up. And now you've got a picture, and nobody trying to crack your computer is - even if your password was weak. And that's what this is protection for. You could, you know, your password could be "Hi Mom," and the bad guys could guess "Hi Mom," and it wouldn't work because they don't have the unique picture which is your second factor that goes along with your weak password in order to strengthen it. So you've got great protection.

MIKE: Great news. Okay. No. 9, a listener in Asia requesting anonymity wondered about corporate man-in-the-middle attacks and EV or extended validation certificates. He writes: Hi, Steve. Today was the day that my employer put an appliance to filter SSL traffic into production without telling anyone beforehand.

Steve: Oops.

MIKE: After one of my development tools with pinned certificates broke, I could quickly verify with your Fingerprints webpage that my employer was spoofing SSL certificates. Your website was also handy to explain to colleagues what was going on. Two things really surprised me, and I would love to hear your opinion. No. 1, I noticed that my employer did not seem to spoof EV certificates, e.g., GRC.com's EV cert was not spoofed, while Google.com's non-EV certificate was spoofed. I'm wondering whether there would be any technical reason why they would not be able to spoof EV certs.

No. 2, while Chrome on Windows seemed to be perfectly happy with spoofed certificates, Firefox was extremely unhappy, thanks to its independent certificate store. This makes me wonder what the reasons might be that Google Chrome tolerates corporate certificate spoofing. I'm also wondering if Google has anything in the pipeline for Chrome to make it easier to detect corporate certificate spoofing or even prevent it. If you pick my question, please don't mention my name. Thanks.

Steve: So I actually, when I was putting the Fingerprints page together, I encountered something that I had not been aware of, or I think I remembered it, but I never really wrestled it to the ground. And that is that another unique feature of extended validation certificates is that a characteristic of them is that browsers have the EV root specifically embedded in them. That is, it is specifically for the purpose of preventing this kind of EV certificate spoofing that browsers do not rely on the standard CA hierarchy. They pin the signing of those certificates to the certificate signers, and they incorporate those certificates. So Firefox is refusing to - well, okay.

So first of all, everybody is refusing EV certificates except Internet Explorer. For a reason that I cannot fathom, Microsoft has made it convenient to set the EV character or characteristics in certificates that people, like in corporate environments, make for themselves. I have no idea why a corporate user - I can understand why they want to make their own certificates for their own servers to use on their Intranet. Why they have any need for, like, essentially fake extended validation, I will never understand. I think this is a horrible mistake because extended validation needs to mean something. And the fact that Microsoft lets people make their own certs and turn that on - it's there in the user interface, click this button if you would like extended validation - is beyond me. So IE is consequently worthless for actually detecting this. But Firefox is now probably the strongest browser for that. Chrome is no longer because Chrome uses Windows built-in certificate store.

So the reason that this listener saw that Google's Chrome did not mind spoofing is that they're in a Windows environment, and a group policy was pushed out through the Intranet, forcing every machine at logon to accept this spoofed certificate from the proxy which is now intercepting SSL communications. That's why nobody had to agree to anything or click anything or say yes, I trust anything. It's possible for Windows systems to acquire certificates through the group policy system.

So suddenly all the systems in the Internet this morning, when booted, acquired this new certificate. And IE trusted it because it's in the Windows root store. Chrome trusted it because Chrome uses Windows' root stores for its certificates. Firefox said, what the heck, I'm not trusting any sites because I don't know what the certificate is. So Firefox you can trust. Safari I believe you can trust. And probably Opera, although I think, now that Opera has switched over to use the Chromium codebase, it's probably now also no longer trustworthy for this. When I wrote that page, it was, but probably no longer. So that's what's going on. Basically EV is an extra level of useful test, but you need a browser that'll complain about it. And I'm afraid that only Firefox provides that assurance now. Unfortunately, Microsoft has allowed their own certificates to have EV flags turned on for reasons, as I said, I just can't fathom.

MIKE: All right. Well, No. 10. John Bailey in Villa Park, Illinois wonders about the relationship between the CA list in browser versus the CA list in the operating system. He writes: Steve, in Security Now! 501 you and Leo talked about the CA lists and how to delete them, and also why it's probably not worth the effort. I looked at the CA list in my Firefox 37.0 browser and found the Hong Kong Post Office, et cetera. I looked at the Certificate Manager utility in my Windows 8.1 Surface Pro 3 and found a modest list of 24 CAs and 40 certificates. How does the browser's list relate to the Windows 8.1 list? By the way, my bank, BMO Harris, looks worse than BofA when you run the SSL Labs test.

Steve: Yup, so there was another listener who, like, looked at his own bank after we brought the attention to BofA, and it's like, okay, how could it be that they are, like, I mean, why banks? You'd think that's where security would count. But maybe not. Okay. So this is really interesting. I talked about, years ago, discovering 400-some CAs, certificate authorities', certificates in my Windows XP root store and thinking, oh, my god, I had no idea we were trusting this many random people, like CNNIC and Hong Kong Post Office. Just, I mean, actually Hong Kong Post Office has been the whipping boy ever since because it's like, what? Why am I trusting anything that the Hong Kong Post Office decides to sign? But, you know, we are. Firefox is using that model, that is, the browser brings the certificate store with it.

What I learned when I installed Server 2008 R2, which is the server I upgraded my servers to the holiday season before last, so Christmas of 2013-2014, I looked at the certificate store. And just as John saw in his Windows 8.1, it was almost empty. And, I mean, this was a brand new installation. And it was, like, maybe it was empty. And it was like, wait a minute. What happened to the certificate store? I did some research. It turns out Microsoft has changed the way they operate. And I don't really understand the logic. But they now have, since whatever, since certainly Windows 2008 R2, an on-demand certificate root store. That is, when your browser goes to a site and gets its certificate, it follows the chain and asks Microsoft on the fly for any certificates not provided in the chain, and the root certificate that by definition cannot be in the chain, and says, "Hey, I'm being asked to trust this certificate. Should I?" If so, on the fly, Microsoft provides that to the system.

So this is actually very elegant. I can't explain it. Because what it means now is, rather than having a ridiculous number of certificates that you're absolutely never going to need, you only get the ones that you need. So John, while using 8.1 for some period of time, only encountered 24 different certificates that he actually needed in his root in order to trust all the sites he's been to since he started using that system. And that's the point that I was making a few weeks ago where I said, boy, you know, there are hundreds of certificates, all of which we're trusting. And the problem is very few are needed for most people. And this paints a perfect real-world example of that.

Now, the problem is you will get from Microsoft any certificate you need. So this isn't actually providing any protection. It might, if it prompted you with a dialogue box saying, hey, you've been using Windows for a year, and you've never asked for this certificate before. Are you sure you want to add it to your store? We're happy to, and up until now we've been doing them all automatically for you, but we thought, yeah, you know, things have settled down after a year. Maybe you want to - now maybe we should ask you. That would be nice. Microsoft doesn't do that. And Microsoft does recognize that getting them on the fly might be a problem. So there is a knowledge base link you can click to get the entire 400-plus blob of CA root certs and dump them all in at once, if for some reason you don't want to get them on the fly. I think getting them on the fly is cool, not to just have hundreds of extra certs that you are absolutely never going to need.

So that explains the mystery, John. That's the relationship. Firefox, being its own store, has to bring them - I guess they could do the same thing, but they don't. They just bring them all along, and they're there. There will be a little bit of a time penalty when you're making that first connection to a site using a root whose CA you don't yet have loaded in your machine because your machine has to go ping the mothership and get that cert and then install it and verify the chain of trust and then say, okay, yeah, fine, we're good to go. But, you know, that's only the first time per CA. And you'll very quickly acquire all the ones you need, as John has demonstrated. But Firefox just brings them all along, as has been the way it operates traditionally.

MIKE: And that was No. 10. My gosh, this is a great show, and it's so great to not only be able to listen to it, be able to participate in it. Normally I just listen while I'm doing the dishes. So this is really fantastic, Steve, and it's such a privilege to co-host the show with you. Any parting thoughts, before we close this thing out?

Steve: No, I just wanted to say it's been great working with you, Mike. It went smoothly, and I think we did another great podcast together. So next time, I'm sure there'll be a next time. I think Leo keeps talking about trips that are coming up. So hopefully we'll be seeing more of you on the podcast. Love to have you back.

MIKE: Yeah, we should send them off to some fabulous vacation somewhere so I can do the show again, absolutely true. And of course tonight is the big podcast awards, and I'm rooting for you. And you should win. And if you don't win, there is just no justice in the world at all, as far as I'm concerned.

Steve: Well, I have a feeling we've stacked the deck because we had a lot of listeners who were saying - and, I mean, it was crazy, too, because for some reason they were, like, vote often, vote early. Vote early, vote often. You could vote every day. It's like, what kind of a voting system has everybody, like, I mean, and they would ask you for your email. I voted for myself, I'm not ashamed to admit. I wanted to win.

And so I used my Gmail account, and I got a confirmation, and I clicked on the link to verify. And then I came back within 24 hours, and they said, oh, it hasn't been 24 hours. So okay, I waited another hour or two, then I voted for myself again. I didn't do it every day. But still it's like, that's just loony tunes. Why wouldn't they just get one vote per person? But it's not the way they wanted it. That's not the way they set it up. So maybe they figured that asking people, giving them the opportunity to vote every day would be a more accurate gauge of people's fervor. I don't know. We'll see what happens.

MIKE: Well, I mean, I just think there's been a lot of acceptance for that kind of thing because of online polls, of course, are just like that. You can vote as often as you like. But if somebody's going to stack the deck, it might as well be your audience. So somebody's got to do it. Anyway...

Steve: So we may talk about the Great Cannon next week, China's Great Cannon. The technology is staggering, and there's some neat diagrams of it that are up now. So unless anything else happens between now and then, I think we're going to go into this interesting sort of escalation of the international Internet arms race and how it works and what it means.

MIKE: Yeah. I'm looking forward to that because I need answers, answers. Leo and Steve do Security Now! at 1:30 p.m. Pacific, 4:30 p.m. Eastern, 2030 UTC every Tuesday, right here on the TWiT network. Of course you can watch live at live.twit.tv. Or you can subscribe at TWiT.tv/sn, or on the podcasting app of your choice, or both, whichever you choose. So this is another exciting episode of Security Now!. Thank you for tuning in. And you will see Leo and Steve back here next week. Thanks for joining us.

Steve: Thanks, Mike.