Listener Feedback #168

Description: Steve and Leo discuss the week's major security events and discuss questions and comments from listeners of previous episodes. They 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-405.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-405-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson's here. We're going to answer some questions from the listeners, talk about the recent security news, and a whole lot more. Will you stay here? Because Security Now! is coming up next.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 405, recorded May 22, 2013: Your questions, Steve's answers, #168.

It's time for Security Now! with Mr. Steve Gibson, the Explainer in Chief, the guy who protects our privacy and security online. Hey, Mr. G. How are you today?

Steve Gibson: Hey, Leo. It's great to be with you again for Episode 405.

Leo: 405, wow.

Steve: You missed the missing Episode 404.

Leo: Did we just skip right over that?

Steve: 404 not found. Actually it would have been fun, although it would have been a constant source of concern for people who are like, I mean, we would have been getting email for the rest of our lives. Hey, where's Episode 404? It's like, okay. So anyway, Iyaz and I had a...

Leo: It's like a hotel that doesn't have a 13th floor, you know?

Steve: Exactly. Iyaz and I had a great time last week, and I know you did at Google I/O.

Leo: I had a blast, and I got me one of them Chromebook Pixels.

Steve: Yup.

Leo: Yeah. That's about all I can say about it.

Steve: So this is a Q&A episode, our #168. And I've got some interesting news. We didn't have a huge news week. The thing that really excited me, I couldn't believe it when I caught up with this news. This apparently was - it's a few months old, so it's just as well that I hadn't mentioned it before because it turned out to be too good to be true, kinda. And that is that our friend Jonathan Mayer, who is not - Jonathan Mayer or Mayer, not the vocalist, but the security guy, the security and privacy guy at Stanford. He's the person who reversed the mistake that Adobe made, essentially. And we've talked about Jonathan through the years because he's been very active in privacy. He reversed the mistake that Adobe, I'm sorry, not Adobe, Apache, the Apache server made on the fact that it was going to be removing the presence of the DNT header from IE10 because somebody at Mozilla who was also, or I don't know if they were at Mozilla, but one of the other people involved in the Apache project just decided they wanted to punish Microsoft for not getting sufficient end-user permission.

Anyway, I installed, and we talked last week about the most recent release of Firefox, v21. So when I saw that Firefox v22 would be by default blocking third-party cookies, I was very excited because that's big. I mean, that's huge. And it may be a bridge too far. It may be more than we can get, but we're getting very close. Right now the only browser that does this is Safari. But that's saying a lot. I mean, Safari blocks third-party cookies by default. So that's a big deal.

So Jonathan blogged in February of this year, so several months back, that this had been submitted into the flow of Firefoxia, which is moving forward, and it would be coming out in version, yes, in v22 of Firefox. He said - and this was on the 22nd of February. He has WebPolicy.org is his blog. And the title of the blog was "The New Firefox Cookie Policy."

He said, "The default Firefox cookie policy will, beginning with release 22, more closely reflect user privacy preferences. This mini-FAQ" - and I'm just going to read two, the first two - "addresses some of the questions that I've received from Mozillans, web developers, and users. First question: How does the new Firefox cookie policy work? Answer: Roughly, only websites that you actually visit can use cookies to track you across the web. More precisely, if content has a first-party origin, nothing changes. Content from a third-party origin only has cookie permissions if its origin already has at least one cookie set." Now, we'll come back to that in a second.

And then, second question: "How does Firefox's new policy compare to the other major browsers?" So his answer was "Chrome: Allows all cookies. Internet Explorer: Cookie permissions vary by [the so-called] P3P compact policy," which is something we've discussed in the past. But continuing, it says, "In practice, almost all third-party tracking

cookies are allowed. And Safari: First-party content has cookie permissions; third-party content only has cookie permissions if the content already has at least one cookie set. In short, the new Firefox policy is a slightly relaxed version of the Safari policy."

So I was very excited. This got into my notes late last week, and I was excited. And then I saw an update, let's see, yesterday, 5/21, also Jonathan's most recent posting said, well, it explained that the CTO of Mozilla had said he wants to wait for one more cycle. Now, I picked up on this because it was all over the place. It was BusinessInsider.com talked about it, AdExchanger.com, I mean, Firefox changing their policy to block third-party cookies is regarded as a big deal. IE has tried a couple times. Some betas of IE did have third-party cookies initially blocked. No one ever seems to make it out of the gate. Something happens.

And so it just must be that somewhere there are major advertising gods that descend on the hapless web developers who are trying to ship a browser like Safari, which in all of its incarnations acts this way, and they get stopped. So I'm holding my breath that this can happen because Jonathan's feeling is that it is necessary to go further than the Do Not Track request and actively decline to accept cookies from sites you don't visit. And that's the way Microsoft has - I mean, sorry. That's the way Apple expresses it and Mozilla expresses it, very cleanly. Do you want to accept cookies from sites you don't visit? And it's easy to say, uh, no, why would I want that?

Leo: The reason the Mozilla's delaying it is because of false positives and false negatives. Do you not buy that as a...

Steve: [Strangled sound] They're - yeah. I mean, maybe. Jonathan's argument is, well, this is the way Safari has always worked, and it doesn't seem to be a problem for the entire - anything anywhere in the Safari ecosystem. Now, there is...

Leo: Well, you might have a problem. Don't you serve images from a separate server than your main server?

Steve: Yeah, but it's got nothing to do with cookies. Cookies are just for tracking.

Leo: Well, sometimes people use cookies in that. That was the example that Mozilla gave, Brendan Eich gave.

Steve: So, okay. There is a subtlety that's interesting, which is that - and this is something I've looked at extensively. Five years ago I wrote that cookie forensics - a set of cookie forensics pages. And we deeply characterized all the different browsers. Not a single one of them was bug free. We were looking for, like, cookie handling bugs. Every single browser had problems. Things like you could sneak a cookie through in the favicon query, and the browser wasn't blocking that, little freaky things like that where, if someone knew about these things, and presumably everyone does, although they don't talk about them, you could arrange to get cookies in.

One browser, and I believe it is only IE, has an - because for a while there were some smart security people at Microsoft. It notices what's called the cookie's "context." That is, it notices if you acquire a cookie in a first-party context. It flags the cookie has having

been acquired in a first-party context. And then it will not send it back in a third-party context. Which is very cool because the problem with the Safari approach and this nascent Firefox approach is that it is still very easy to sneak - it's very easy for third parties to sneak tracking cookies in. All they have...

Leo: But, see, that's the other question, is if you do this, doesn't it just force them underground and to use other methods to do the same thing?

Steve: Perhaps.

Leo: At least with a cookie I know it's being set, and I could check it.

Steve: My feeling is I think Jonathan is right. People don't want tracking. And it is only by virtue of the default setting that third-party cookies are enabled. And if they were just disabled, the world wouldn't collapse. The advertising ecosystem would not die. Everything would be just fine. And there would be less tracking and profiling on the Internet. Obviously it's something I have no...

Leo: There would be less analytics, and a lot of the free content that you enjoy would go away. But that's okay.

Steve: I don't think any of it, no. My point is...

Leo: Doesn't that kill Google Analytics?

Steve: Nothing would change.

Leo: Does that kill Google Analytics?

Steve: No, because Google Analytics is hosted on the pages and is making queries itself. And it's running script in your browsers.

Leo: JavaScript, right.

Steve: Yeah. So it's got a grip on everything going on.

Leo: Yeah. I mean, if it doesn't break that kind of analytics, it's not as much of a problem.

Steve: It's not going to break anything. None of this breaks anything. This is just advertisers not wanting any pushback at all. And they're obviously very powerful since

only Apple seems to have pulled this off. We'll see whether the Mozilla folks are able to do it. It would be wonderful if they did. Remember that there's a - I've got a graphic on my site, on the browser statistics page of the cookie pages, that shows, from all the visitors, we have about 66,000 unique visitors a week, and it shows across the browsers who has third-party cookies enabled. And the bar for Safari is completely - stands alone compared to all the other bars, just because that's the default. And it would be wonderful if Firefox had the guts to pull this off, did, and then maybe Opera could follow suit, and then finally IE would be the odd man out. It'd be really interesting to see if Google ever makes this change because of course advertising is their entire revenue model.

Now, last week, I don't know if you picked up on this, Leo. I was wishing that I had you and Iyaz both, just because this was interesting. But you may have seen it go by because there was a lot of attention given to this. It was discovered - and in fact Ars Technica picked up on it days later, which gave it several more news cycles. It was discovered that links people send each other in Skype chat messages are visited by Microsoft's servers.

So the Heise Security that we've spoken of often was the first official security company to pick up on this. They deliberately set up some links, even including logon material in the URL, and used Skype chat to send this to someone and then monitored their logs. And sure enough, a few hours later, a server in Redmond with an IP allocated to Microsoft's network visited those links. Ars Technica picked up on it, duplicated the experiment, and verified it. And there was of course lots, a whole flurry of back-and-forth from this because people were upset that Microsoft was following links that were being sent in Skype messages.

And the point that I intended to make last week when we talked about this was, well, and there were people saying, well, you know, Microsoft is doing this to remove spam from Skype and to control the content. Their EULA states that they have the right to scan messages and so forth. And so my takeaway was just to note that they are able to do this. I mean, it's not like they could or they might, it's that they are. So I just wanted to come back and dot that "I" and cross the "T" that it's not why they're doing it or what they're doing with it, but just absolute evidence that Skype chat text messages are being read in Redmond. It's not like they're being kept encrypted, or they will only decrypt them if the government requires them to and so forth, but that they can.

Which is a segue into next week's podcast because there is a very nice-looking, very secure chat technology, chat system, called Cryptocat: C-r-y-p-t-o dot c-a-t. And so they've published everything, enough that I can do a full analysis of its security. So that will be the topic for next week because, in response to last week's discussion of this, there was just a bunch of email in the mailbag and through Twitter saying, okay, is there a secure chat solution? And I believe there is. We'll have the full readout about Cryptocat next week.

Leo: But this doesn't necessarily speak to voice calling on Skype. It's the messaging, the text messaging.

Steve: Correct. Correct.

Leo: And it's unclear where this all happens. It could be, it probably is part of Microsoft's attempt to prevent the spamming of malware through links. So it's

probably, what they think it's likely doing is it's checking it against Microsoft's Site Advisor database.

Steve: Yeah. Although it is doing it with several hours of delay. So...

Leo: Well, there's a lot of people doing it.

Steve: I was going to say, if there's a backlog...

Leo: The other issue is it's not clear where the request is coming from. It could be coming from, for instance, the receiving client, that Skype could be then unencrypting, looking at it and saying, okay, I see - and in fact this seems like a likely scenario. I see five links in this message. Let me run a scan against these links against Microsoft's Site Advisor to make sure they're not malware.

Steve: Yeah. That came up. And unfortunately I don't remember what the argument was. I think that somebody checked for that and ruled that out.

Leo: I think we don't know. All they know is there was a hit. They don't - they created a phony page, and they got a hit from Redmond. And the only way that that link was disseminated was through obviously not encrypted Skype transmission. I think the real issue is the larger issue, which is that it's clear then that Microsoft could in theory see everything. Maybe they're only checking those links for malware. There might be, for instance, a regular expression parser and a Skype engine, recipient engine that says I see a URL here, let me check it against Site Advisor.

Steve: Right.

Leo: So I just don't want people, I mean, the real issue is the potential that Microsoft could be spying on you. And as a lot of people pointed out, hey, that's what happens with closed search software. If you think about it, any client you're running on your computer, if you don't see the source code, could be doing anything.

Steve: Well, and that's why a solution like Cryptocat, which is open source and multiplatform, I mean, for example, I'm still waiting to hear back from the BitTorrent guys on their protocol. We have a couple - we have a question or two in this week's Q&A that are addressing questions about BitTorrent Sync. And it's like, well, we've gone as far as we can at this point, until we get a readout. Everything looks good. But I need the details in order to see whether I can say yes, given what we know, they absolutely did everything right.

Also there's been all this brouhaha about the Department of Justice and the Associated Press. And the first stories that I read said that the DOJ had the conversations that the Associated Press reporters were having. And later stories corrected that. So I wanted to

update our coverage and say, yes, that's correct, all they got, what the Department of Justice got, and it was their confession, I mean, they wrote a letter to the AP saying, by the way, we have all of this. Way after the fact, but they did acknowledge that, that it was phone call records, that is, who the AP reporters called, not the content of their conversations. So I just wanted to correct that.

And then I looked at this, and I thought - I just shook my head. I thought, well, this is not going to be a surprise to any of our listeners. So I called this the "Uh-huh" news of the week. And this was picked up - this was in BBC Business News. The title was "Contactless 'Charging Errors' at Marks and Spencer," which is a chain of stores in the U.K. And so the story is some Marks and Spencer customers have told the BBC of cases where the chain's contactless payment terminals have taken money from cards other than the ones intended for payment. So these are near field communications cards.

And so the story, to give our listeners a sense for this, it says, "Cards are supposed to be within about four centimeters of the front of the contactless terminal to work. But some customers say payments have been taken from cards while in purses and wallets at much greater distances. M&S, which is Marks and Spencer, said its systems had been extensively tested and were robust. Marks and Spencer recently rolled out the contactless payment system to 644 U.K. stores. The system uses something called near field communication to identify a card and take payment.

"Rosemary from Sussex got a shock when she tried to pay by chip-and-pin at her local store. She believes her contactless Smile card was much more than four centimeters away from the terminal when she visited Marks and Spencer in Chichester in April and tried to pay with her regular Lloyd's debit card. She told BBC Radio 4's Money Box program, 'I put my card into the reader, and the assistant was asking whether or not I wanted cash back.' She said, 'Before I could answer, the transaction came up as complete, and the till issued a receipt. So I hadn't put in a PIN at all at that stage. I queried it with an assistant, and she looked rather puzzled and looked at the receipt and compared it to my card and realized that the numbers didn't tally.'

"Rosemary, however, recognized that the four digits on the till receipt belonged to a different Smile card she had in her purse, which she was holding in her other hand. She had not realized until then that this card was able to make contactless payments. Even when she realized it could, she thought her purse was about a foot or more away from the terminal when the payment was taken."

So anyway, this story continues, talking about other M&S customers who found double payments on their bills from Marks and Spencer where somehow a card that they had not intended to use purchased the item, then they purchased it again with the card they had intended to use. So this is the fundamental problem with radio is that it's just not precise enough. In order for the receiver to be robust and sensitive enough to pick up the card with sufficient reliability in any orientation, it's going to have to be sufficiently sensitive to be able to pick up the card in an ideal orientation further away. And it's just a bad idea. This whole notion of, oh, the magic of it just being nearby and working is crazy.

Leo: Well, and that's why they limit the transactions to 20 pounds, and they require a PIN. I'm not sure what happened with the PIN on this one. Maybe somebody saw the PIN.

Steve: Yeah, apparently - yeah. Apparently didn't need it.

Leo: Yeah, but it does, though. Somebody saw the PIN, probably. I would guess. Not always. So Craig B. says that does not always require a PIN. I thought that was the general idea, but maybe not.

Steve: I think, I mean, again, there's this aspect of glamour where they want it to be magical, where you just sort of wave the card by the terminal, and it's like, oh, the money jumps across.

Leo: That's why you want a PIN. That's why you'd like to have a PIN.

Steve: Exactly. Also another very highly tweeted note came from the runner-up, essentially, in Intel's International Science and Engineering Fair. And the only thing I could think is this is another example of a really good PR firm at work. Because this thing was just - it got an amazing amount of coverage for really a non-event. Now, I guess the event is that the woman who or the young woman...

Leo: It's a girl, that's why, a young girl.

Steve: Okay. And she's 18, yes.

Leo: Yeah, yeah. And so that's cool.

Steve: So she's 18, and so she did something working in a nanotechnology lab where she produced a supercapacitor. And, I mean, I'm not taking anything away from that effort at all. That's not what I mean. But the headlines were "18 year old's invention can recharge a cell phone in 30 seconds."

Leo: Right. I saw that, too.

Steve: Which is nonsense.

Leo: Yeah.

Steve: And it's like, okay, no.

Leo: Well, what if you built the supercapacitor into the phone and then charged it and then had it trickle-charge the lithium ion or something like that? That's what - I got the impression she was putting the supercapacitor - I don't know. I didn't read it very closely.

Steve: No, I mean, this was just...

Leo: We've talked about this before.

Steve: Yeah. All this was, was this was a story about someone created a supercapacitor.

Leo: Except that she didn't.

Steve: And it's like, okay. And she's an 18-year-old girl.

Leo: Right.

Steve: And so I think that's great. But there was a sense of this was actually revolutionary, and it came from an 18-year-old girl. And it is actually, as far as we know, not revolutionary.

Leo: You've been talking about - you talked about this two years ago.

Steve: I know. It's a nice supercapacitor. So...

Leo: All right. So she didn't do anything new.

Steve: No. She's just young, and so...

Leo: Well, it's a good story.

Steve: Yeah.

Leo: And you know how sometimes the press doesn't really know that this isn't new.

Steve: No. And so what happened is our listeners tweeted me like crazy.

Leo: Because we know you're interested in supercapacitors.

Steve: Absolutely. And so thank you for that, and just so everyone knows, we now have another supercapacitor in the world. Okay.

Leo: I have my supercapacitor screwdriver that charges very, very quickly.

Steve: Oh, I mean, yeah. And I'm not buying an eCar because I'm waiting. I mean, we're seeing - we talked about it in the last couple weeks. We're seeing tremendous effort now expended in electrochemical storage and electrostatic storage. These are going to move forward. And the nature of a car is that you can't just swap out the energy source with a better one. It'd be like, you know, is your car diesel or not? I mean, you've got to give it the proper kind of fuel. And you just can't change your mind. So anyway, I'm waiting because I just think we're still on - we're on the cusp of a real breakthrough in some sort of car, you know, automotive-compatible energy storage.

Leo: Right.

Steve: And that's good. I did see a weird note about - I don't remember what country it was. Somebody was talking about storing excess energy by pumping the air out of a container at the bottom of the ocean. And the idea would be they would, like, pull the water out of the container, and then the ocean's pressure to get back in the container would run turbines. And it's like, well, okay, there's an innovative...

Leo: But that's a battery. That's not a creator of energy, that's a way of storing energy.

Steve: Correct.

Leo: Because you have to pump the energy - to pump the stuff out uses the energy.

Steve: Exactly.

Leo: There's lots of - this goes back to the Greeks.

Steve: Yeah, and, well, in fact, there have been efforts to, like, pressurize, to store air under pressure in the ground and then have it come back out. It's like, okay. I don't know how you plug all the leaks. You know, gophers are going to have an interesting life.

Leo: If you think about it, if you hand-cranked water up into a tank, you're storing the energy that you used to hand-crank it up, and then at a later time you can release that energy by releasing the water and having it turn a little water wheel. And there you go. You're invented the battery.

Steve: Hydroelectric dams. How do you think those function?

Leo: There you go.

Steve: There we go.

Leo: There you go.

Steve: So AskMrWizard.com is still cranking away. He wanted me to note that he finished Episode 26, 13 videos for Episode 26. That was one of our classic How the Internet Works series. So for those of you who have been enjoying the AskMrWizard.com series, if you go to AskMrWizard.com/securitynow, that'll take you there, and he's got a bunch of links. So he's got both Episode 25 and 26, and I assume he's working on 27 because it was, I believe, a three-part series where we covered sort of, like, the entire fundamental architecture from packets to routing to ports and everything, about how the 'Net works. And he's animated graphics to go along with it.

Leo: And he puts ads. Just so people understand, he puts ads in it. This is why he's doing this.

Steve: Oh. I didn't know that. Okay.

Leo: Yeah. He's making a little money on it. That's all right. That's completely legitimate.

Steve: Yep. So do we.

Leo: Yeah.

Steve: And I did tweet a project that I thought was very nice. So it's a Kickstarter project. It's called "meta," m-e-t-a, "the most advanced augmented reality interface." And I don't know if meta is an acronym. Maybe most, then - I don't know. I don't see how you can get META out of the Most Advanced Augmented Reality Interface. But they are very nice-looking glasses, and they've put together a compelling video, and they've got a bunch of smart augmented reality people on the team. The idea being you've got essentially what looks like a stereo camera looking down in front of you. And what's so cool is you can use your hands to manipulate stuff in the air in front of you. And there's one great picture on there that shows a person putting his hand out, and each of his fingers is circled with a different color. And then in the video he's, like, moving things around. And they show a bunch of people messing around with some architecture, changing the planting of palm trees by reaching out into space and moving things around.

So anyway, I thought it's just cute and cool. And I tweeted it, so people who follow me already know about that. But those listeners who don't, and there are many more of you, I wanted to sort of bring that to your attention [kck.st/12rTG3J].

Now, Leo, Star Trek.

Leo: You went to see it.

Steve: Oh, my goodness, yes. Opening day.

Leo: Oh, good.

Steve: Now, what I'm seeing is - I've looked at all the reviews. I've read all of the blog posting comments. And I completely, completely get it, that there are purists who will never, apparently, forgive J.J. Abrams for not being Jonathan Frakes. And they will not forgive what's his name, our new Kirk, I can't think of his name...

Leo: I can't remember his name, either.

Steve: Yeah.

Leo: He's good. I like him.

Steve: Yes, he is. Well, for not being Jean-Luc Picard, for not being Patrick Stewart. There were, like, the thinking man's...

Leo: Chris Pine is it, yeah.

Steve: Chris Pine, yes. There were the thinking man's Star Trek movies. And...

Leo: I like them as the young, I mean, that's what's cool about it.

Steve: Oh, Leo.

Leo: The young Spock and the young - it's like watching the, what is the Looney Tunes that has the little kids? Animaniacs? It's the Animaniacs Star Trek.

Steve: I love these. I mean, and I tweeted...

Leo: Tiny Toons, there it is.

Steve: I was exhausted after two hours of watching this. I mean, it was fantastic. I loved this new movie. So, I mean, it is an action sci-fi adventure movie.

Leo: No spoilers.

Steve: No, I'm not, I'm being very careful not to say anything.

Leo: No, you would never do that, yeah.

Steve: No, I would never do that. And for those people who see it, I will say I was a little disquieted by some of the decisions that were made. But...

Leo: Oh, there's some massive plot holes. But that's okay.

Steve: Yeah, yeah.

Leo: But, you know, what do you expect?

Steve: It was a great movie.

Leo: It was fun to - it was exciting. And what you can't get over, and I think this is - to me, this is the whole appeal of all the Star Trek movies, going back to the first one, is kind of the thrill of seeing these characters we love again.

Steve: Yes.

Leo: And that, you know, I remember the first time Kirk and Spock, the originals, Shatner and Nimoy, walked onto the big screen because it was many years after the TV show.

Steve: Oh, oh.

Leo: And it was just like, "Oh! Old friend!"

Steve: [Laughing] Yes.

Leo: And this has that same kind of thing. And they're smart, J.J.'s smart enough to throw in a lot of references to older stuff.

Steve: Oh, you're right. There's plenty of that for us.

Leo: That's what makes it fun, yeah.

Steve: And the point I made to some people who I was talking to was, look, this is for making money. I mean, this is to make...

Leo: What? No. They make money?

Steve: Fundamentally, this is about money.

Leo: That's what tracking cookies are for. They've got to stop this!

Steve: [Laughing] And so he's creating a popular movie with broad appeal. A good friend of mine took his mom, and his mom loved it, and now she wants to see the first one and also one of the older original movies that shall not be mentioned. So, and I must say, I give credit to somebody because there were some things I read in anticipation of this that were clearly deliberately misleading about the plot. And I thought, oh, good for them. They had me completely going off in the wrong direction, thinking that I knew something about it, and I didn't. So that was good.

Leo: You know who I liked best? I shouldn't say "best," but who I really am enjoying is Bones in this.

Steve: Yeah, he's good.

Leo: He's really good. And of course Scotty is fabulous.

Steve: And Bones is...

Leo: Bones is Bones. "I'm a doctor, Jim, not a...."

Steve: Yes. "Goddammit, I'm a doctor, not a..." whatever.

Leo: We won't say it. I don't want no spoilers.

Steve: Yes.

Leo: "I'm a doctor, Jim."

Steve: Okay, now, I know you're seated. You're on your inflatable ball?

Leo: I'm on my ball, yes.

Steve: Okay. Because this is so painful, Leo.

Leo: Oh, no.

Steve: But I really do...

Leo: Well, it's good I'm on my ball, then.

Steve: This is so painful. This is thanks to our friend Simon Zerafa, who somehow finds what we would call humor. This is just awful. I mean, this is so awful I had to share it.

Leo: Oh, great. Thanks.

Steve: An SQL query goes into a bar.

Leo: Oh, no. Yes. I like it already.

Steve: So we have a SQL query goes into a bar, walks up to two tables and asks, "Can I join you?"

Leo: Oh, that's pretty good. You know what? I like it. Can I - goes to two tables and says, "Can I join you?"

Steve: Can I join you. So for those who don't get it, it's like, oh, that's okay, you're not supposed to. For those who do, ouch, you know, I apologize.

Leo: I still think that the Little Bobby Tables xkcd comic book is the best...

Steve: Yes, that was really...

Leo: ...MySQL joke I've ever seen.

Steve: It's brilliant.

Leo: I'll just show it, for those who haven't seen it, it's xkcd No. 327: "Hi. This is your son's school. We're having some computer trouble." "Oh, dear. Did he break something?" "In a way. Did you really name your son Robert '); DROP TABLE Students;--?" "Oh, yes. Little Bobby Tables, we call him." "Well, we've lost this year's student records. I hope you're happy." "And I hope you've learned to sanitize your database inputs." [Laughing] Robert'); DROP TABLE Students;--?

Steve: Very clever.

Leo: [Laughing]

Steve: Very clever.

Leo: There's something about MySQL jokes that just crack me up.

Steve: Okay. So rarely do we ever hear any complaints from people that SpinRite did not spend long enough on their drive.

Leo: No, I've never heard that.

Steve: No. So Charles in Sydney asked me, he says, "Hi, Steve. I refer to a recent podcast where a SpinRite user said it took SpinRite four weeks to recover his mother's drive. I recently ran SpinRite on a 60GB PS3 disk which was having problems, and it only took one hour to run and told me that I had one unrecoverable sector. Why didn't SpinRite take longer to try to recover this sector? I assume it was because it somehow knew that, no matter how hard it tried, it wouldn't be able to recover it. Is that right?" And he said then, "By the way, SpinRite fixed all the problems..."

Leo: Well, that's why.

Steve: "...I was having with the PS3 disk, and everything works perfectly now."

Leo: Oh, okay. Horrible. What a fate.

Steve: Charles, I'm sorry that it didn't make it feel like it was working harder to recover your drive. I actually saw a slowdown in sales after that four-week story because I think people were like, wait a minute. This thing takes four weeks to recover a drive?

Leo: Oh, no, no, no.

Steve: It's like, no, no, no. No. I mean, it will - and I once said it will take as long as you want it to. But in this case it didn't need any more time, so it didn't take any more time. So Charles was right. SpinRite did figure out that it was - it did everything it was possible to do on that sector, fixed his drive, and only took an hour. And it only typically takes a few hours. So, no, no. Not a few weeks. A few hours.

Leo: It's related to how hard a sector might be to recover and how many bad sectors there are.

Steve: Actually, one of the - yes. One of the problems is that, if a transfer ever creates an error, the formal specification says I have to tell the BIOS in the motherboard to reset itself. Some BIOSes take a long time to do that. So that process slows things down. So Charles had a motherboard that was quick to reset, and so it didn't take SpinRite that long to get through it. And it's one of the things that I will definitely be looking at in the future, seeing if I can speed it up for everybody by not using the BIOS any longer, which will - everyone will breathe a sigh of relief. But for what it's worth, typically it doesn't take four weeks. It typically takes hours.

Leo: Do you work on UEFI computers? Because now all the new Windows 8 machines I think are UEFI, not BIOS.

Steve: Yeah, we do. We've never - because UEFI boards are out there. They all have a BIOS emulation layer.

Leo: Anyway, right.

Steve: Yes. So they can...

Leo: So you can do an N13 or whatever it is you do.

Steve: Yep. We still work just fine.

Leo: All righty, Stevie. Got some questions for you, Steverino. This comes to us from Dorset. And he wants to make sure we pronounce it right, so Dorset, England. Andrew Stevenson. He says that apparently you've come to the attention of the NSA. Oh, how exciting for you, Steve.

Steve: Oh, you, too, Leo.

Leo: Me, too? Oh, crap.

Steve: Oh, yeah.

Leo: I just wanted to drop you a line and say that a recently declassified NSA document has appeared online that directly links - oh, god - directly links to GRC and features ShieldsUP! and Security Now!. The document also names both of you. Oh, double crap.

I was reading a Wired article about the tips NSA has on spying using Google - that's on the Wired magazine's Threat Level blog, which is really great reading. Linked from the article is the following 640-page, 40MB PDF, "Untangling the Web," from the NSA. If you go to page 583 - hard-coded page, not the PDF reader page - you'll

see a graphic of ShieldsUP!. Page 605 also references the Security Now! podcast. So we've got a shortened link: goo.gl/fNDka. The author of this document has been redacted. But we know he's a fan. Publication date also unclear, although February 2007 was the date of last modification. Well, that's cool. I really like that.

Steve: It actually - so, okay, a couple things.

Leo: Did you go read it?

Steve: Oh, yeah. So here I'm holding up a picture.

Leo: "Untangling the Web: A Guide to Internet Research."

Steve: Yes. And I really would commend this. First of all, 40MB, it's a huge thing because it's scanned. So it's scanned images and compressed PDF.

Leo: Lots of paper, yeah.

Steve: Yes. But apparently it was classified. Now it's got - and it used to say "For Official Use Only," and that's now been crossed out of the entire document. And it's got a Doc ID 4046925. But the thing that was most flattering is there's a chapter, the second reference down on the documents page 605, which is not the physical page but the actual number on the page itself. It's a chapter titled "General Security and Privacy Resources."

And it reads, "The best defenses against the many dangers lurking on the Internet are awareness and information. Because security and privacy threats are so pervasive and increasing in number and potency, staying on top of threats and means of protection is crucial. Steve Gibson, rightly famous for his ShieldsUP! website and free software," and then it says, "e.g., Unplug n' Pray, launched a new service with Tech TV's Leo Laporte in 2005. Every Thursday afternoon," and originally we were Thursday afternoons, "they create a 20- to 25-minute," and originally we were less than half an hour, "audio column about personal computer security called Security Now!. The topics covered include Personal Passwords," he says, "(a must-read), NAT Routers as Firewalls (another must-read), HoneyMonkeys (no, I'm not making that up)," he has in parentheses, "Unbreakable WiFi Security, and Bad WiFi Security.

"The audio broadcasts are archived in several formats, including a text file, a PDF version, and an HTML web page. There's also an option to receive an email reminder whenever the page is updated. Gibson has the ability to cut through the jargon to explain these topics clearly and to offer practical advice on how to handle personal computer security issues." And then they have a link, Security Now!, GRC.com/securitynow.htm. So...

Leo: Wow.

Steve: Wow. That was very cool.

Leo: That is really neat.

Steve: And for what it's worth, this is, I mean, it is a fabulous resource for, I mean, as many people have noted, this stuff tends not to age. The stuff we were talking about at the beginning of the podcast, Episode 1, 2, 3, 4, 5, I mean, HoneyMonkeys maybe a little outdated now, but it was a fun podcast anyway. But the essence of what it takes to understand security is timeless. And this is a great PDF for people to scroll through. And clearly the person who brought this to my attention, Andrew Stevenson - thank you, Andrew - had to have gone through it page by page, and he stumbled on references to our work here.

Leo: That's really neat.

Steve: So that was very, very cool.

Leo: Well, and now that we've come to the attention of - yeah?

Steve: It was funny, too, like, I tweeted this to my followers. And I got back a lot of, like, oh, well, very much like your reaction, Leo, when it's like, uh, you've come to the attention of the NSA.

Leo: We're in a database somewhere.

Steve: That's a good thing?

Leo: They're watching now. So, hi. And maybe they're learning some things, which is good, too. Next from Matt, also in England, in London - that's pronounced lawn-dawn - says, Steve, you forgot to secure your cookies. So he's got a long little clip here.

Steve: Well, yeah. What he did was - first of all, he's right. I made a big point a couple weeks ago of talking how I went to HTTPS exclusively. One attribute which cookies can receive is called "secure." So in the set cookie header which the - the so-called "reply header," the reply meaning that the server is replying to the client's query. In the reply header you specify the cookie's name and then its value. You can also specify the path, which is to say for all future queries that are subsets of this path, that is, within this path, return this cookie. But you can also say "secure=1." And what that says is only send this cookie in the event that you have security established. That is only over an HTTPS, which is an SSL/TLS connection.

And he's right. I have not gone back into the source of my so-called "net engine" code and added that everywhere. And I should. So as it turns out, nothing I'm doing requires secure cookies. My entire eCommerce system is cookie free. It operates with no cookies

enabled. And I can maintain session status and so forth without cookies. So the only reason I have any cookies was for that analysis system which allows me to generate statistics because I was curious to see how people had their browsers set and how different browsers were acting and so forth. And I've just let it run for the last five years because I already had the technology in place.

But Matt, you're right. He also noted, or in these reply headers, somebody had suggested I add - and I think it was in our newsgroups - an X-Frame-Options header, which is very cool. One of the things that can happen to a website is that it can be loaded into another site's frame. So some other site would be actually hosting GRC in a frame. Essentially, it's sort of a way of stealing our content. And there are security implications, obviously, to that. I want someone to come to GRC in order to get GRC's stuff, rather than run us in a frame and perhaps play some games with users.

You can actually have your server send a response header that's X-Frame-Options, and mine is set to Same Origin. And essentially that prevents GRC from being framed in somebody else's site. So the browser sees that and will not - it will refuse to - and all the browsers support this now for a long time. The browsers will refuse to present content from a different origin if the server says only show this in the same origin. So it's an additional way of enhancing security, which I'm glad to do. And I'm going to go back and add - turn security on on my cookies because there's no reason not to have that turned on. So thanks for the heads-up, Matt. And everybody should.

Leo: Third question from London.

Steve: Wow.

Leo: Richard King in London had his alarm bells go off. Hello, Steve and Leo. Long-time British viewer. That's in England. The InfoSec Europe Expo has just been held here in London, and I've received loads of freebie updates. Here's one from WickHill.com. It's their SSL Intercept demo. He said, it made my alarm bells go off. I completely follow your Trust No One - I'm just making people - all of our U.K. followers are just hating this right now. I completely follow your Trust No One, and I think I understand your teachings on SSL/TLS; that is, my PC checks the SSL cert, its authenticity, et cetera.

The thing is, for this SSL Intercept - this thing is for SSL Intercept, and it says it's a couple of devices that intercept the SSL/TLS stream, decrypt it for reading, then send it out via a second device which re-encrypts the stream. It doesn't say it uses its own SSL Cert, nor that you have to install it into the browser of every PC as a Trusted Cert. And it says the process works in both directions.

Now, in your recent episode you explained how IE can be tricked into turning green and cause us to trust a spoofed EV cert. Maybe now it's just sinking in. Shouldn't we be paranoid about this thing? How can this be trustworthy? Surely this is a man-in-the-middle scenario? Imagine I am at work or somewhere, and I have not been informed that this is in place. I try to buy my copy of SpinRite with my Visa card. And while I'm imagining the Yabba-Dabba Doo sounding in your office, this device is decrypting my Visa details and reading them. Hmm. If this is the case, then the SSL/TLS model or its implementation is broken, and we need something new. Or have I got this wrong? Please tell me I've got it wrong, Steve. Can you explain,

please? Many regards, and keep up the excellent work, both GRC and TWIT. Richard King, near London, England.

Steve: Okay. So...

Leo: I imagine you clicked this link.

Steve: Oh, yes. And there's a nice, yeah, you can click the link, and it shows you a product which is being offered. It's exactly as he explains. There's one device that does the decryption. And from their diagram it looks like they're not trying to be in the antivirus, spam, and content scanning business. They allow an intermediate device to do that. And then you go back to another one of their devices, which reverses the process, exactly as Richard says, and reencrypts is so that it goes off to the Internet.

So I don't mean to beat a dead horse. I know we've talked about this a lot. But it is central to the security of the Internet is SSL, the transport layer security, HTTPS and what that means. So there is one fact that Richard didn't address. He mentions that they don't say that your PC checks their cert or how authenticity is verified. The SSL ecosystem, the so-called PKI, the public-key infrastructure, is not as broken as this makes it seem because they have left out the fact that your browser must first trust this vendor to issue any certificate it pleases.

Well, that's the horrible breach. That's the problem, is if they can get a certificate into your computer that allows them to synthesize certificates on the fly, and your browser to trust those. So that's the gotcha. That's why this is generally - this is not going to happen if he's out roaming around in a coffee house or outside of his corporate environment because the IT department first has to put a certificate into every employee's computer, which can happen transparently when they log in, for example, to a - Microsoft has this technology, Active Directory, and also Group Policy systems, which are enterprise network management technologies. So Microsoft can just slip a certificate in without you knowing the wiser because that's what their technology allows, making it appear transparent. But if you were somewhere else - and that requires that your computer be configured and log into a so-called domain controller within a Microsoft network in order for all this to happen.

But in general, if you had a noncorporate machine, if you just brought a laptop in from home, hooked it up to the network, and tried to go outside over a secure connection, you would immediately get a notice, like I'm sure we've all seen before, where your browser says, whoa, this site is not trusted.

And in fact I think under Google, if you do that on Chrome, it says something fun, like, "Get me out of here," and it comes up all red. And it's like, there's just no - it's very hard to push yourself through that warning because Google makes it very clear that you don't want to proceed. Firefox brings up something similar, where it's like, I know what I'm doing. I want to go anyway. But so it is that kind of notice that you would normally get unless there's been somehow some preexisting modification to your laptop.

So this is the system that we're using. It's the best anyone has come up with. The extended validation, I don't know if I've mentioned yet that there's a new page on GRC that I finished a couple weeks ago specifically discussing exactly how the enhanced validation EV certificate technology works, and why we have confirmed that you can trust

that in Firefox and Chrome. I know that I talked about this before. The page is finished. It's up in GRC's main menu, and it's easy to find. It's under the - I think under the Research tab. So it's not as bad as Richard worried. You will get a notification unless you have essentially turned the configuration of your machine over to somebody who controls it. And then, obviously, all bets are off.

Leo: And I would presume this is really for business, right, to monitor employee traffic, that kind of thing.

Steve: Yeah. And there have been...

Leo: You know they do that. I mean, this has been around for years.

Steve: Yes. Yes. I've had some pushback from corporate IT who said, gee, Steve, you're painting this in a real dark fashion. I mean, this is just so that bad stuff doesn't come in over secure connections. We want to be able to scan our network to protect our employees. No one is looking at their mail. Nobody is looking at their Visa credit card payments. And I completely agree. My position is users should know that their credit card information is being decrypted and scanned. That's all I want. I just want users to understand that. They can decide, then, if they would rather wait till they go home to purchase something online because...

Leo: I wish they would. Please don't be shopping on my company time.

Steve: Exactly. Exactly. So that's all.

Leo: That's my job. I do the shopping on company time.

Steve: Exactly.

Leo: Yeah. And we, you know, the courts have upheld this. We've mentioned this many times. You're using company resources. According to the law, which by the way differs for analog traffic to digital traffic. The law isn't quite fair in that regard. But from digital traffic the law has always said, and courts have upheld, that employers have every right to monitor what you do.

Steve: Yeah. And we just - I'm all about empowering the end-user. And so this is a way for people to verify.

Leo: Yeah. And there's a lot of good reasons why a company needs to do that, not just malware. There's all sorts of good reasons. And if a company does do that, they should really have an Internet policy that states that clearly, I think. But they're not obligated to.

Steve: No.

Leo: Marsh in Sacramento wonders if Chrome has changed something. Hmm. Long-time listener, love the show. Steve, I know you're a Google Chrome user, and the last few months there's been a change to the way Chrome works for me. I wonder if you've seen the same change because it sounds like the kind of thing you'd scream about. A few weeks ago, random words on web pages started highlighting themselves. These become links to ads of various types. It's very annoying. Also, sometimes I click a link on a page; and, in addition to going to the link, a new browser tab opens, taking me to a survey page. Is this a Chrome thing, or have I been infected by some adware? Thanks, Marsh.

Steve: So what I know is that it's not a Chrome thing.

Leo: It's not Chrome.

Steve: And I have seen sites, not specifically under Chrome, but in probably Firefox...

Leo: All browsers. It works on all browsers. This is a crappy advertising model.

Steve: Yes. And I knew that you would know about it, so I turn this over to you, Leo.

Leo: Well, it's just I hate it. When I see this on pages, I never go back.

Steve: Right.

Leo: But it's a way of embedding ads. You'll highlight words kind of at random. And the thing is, it's deceptive, in my opinion, because it looks like a - it breaks the Internet, the agreement about what the Internet does. It looks like it's a hyperlink, and when your mouse hovers over it, you get a pop-up ad. Whenever I go to a page that does that, I just leave immediately. I think it's disgusting.

Steve: And so it must be that the site...

Leo: JavaScript on the site.

Steve: Yeah. So the site has chosen this approach to monetizing itself.

Leo: Right.

Steve: And I agree, it's really intrusive.

Leo: Just vote with your feet. I presume that the survey is exactly the same kind of thing. And JavaScript can do this. This is a, you know, on mouseover, pop up an ad. That's an easy thing to do. There are a lot of sites that do it. I wish they wouldn't. I understand they need to monetize. But I think that this is a very intrusive and unpleasant way to do it.

Steve: So my guess is that Marsh has just been, by happenstance, choosing sites - he says "run across sites" - that are doing that. And so as you say, go somewhere else.

Leo: Yeah. And NotScripts will stop this. My attitude is it's better to go somewhere else than to respond, retaliate by consuming their content and turning off their monetization. That's the deal with the devil they made. So you go somewhere else and read the content. It's never - it's yet to have been on a site that I have to read, I can tell you that right now.

Steve: Yeah, exactly. It's some link you happened to choose through Google. And it's like, oh, okay, I'm not going to in there in any longer.

Leo: The other potential possibility is that it is some malware on his system. And that could very well be.

Steve: Absolutely.

Leo: NoScript, or NotScripts in Chrome would block that. And I would check for malware because that sounds like the kind of thing malware might do.

Steve: If suddenly - okay. If you go to GRC, and you see any of that nonsense...

Leo: Exactly.

Steve: Then you know, yes.

Leo: If you're on TWiT and you see it, then you know that it's something in your computer, not Chrome, that is infecting your system. That's a good test. Neither you nor I do that.

Steve: Oh, my goodness. And that would never happen.

Leo: I wouldn't dream of it, yeah.

Steve: Unh-unh.

Leo: Dan in Canada Dan in Canada offers a terminology correction: Steve, just an FYI. You know, I knew when you did this, and the chatroom mentioned it. It's not important.

Steve: I know. But I figured, what the hell.

Leo: It's worth a correction. In modern Linux systems, the userland tool to control the firewall is iptables. Ipchains is what you mentioned, was Linux 2.4. So in 2.2 it was ipfwadm. Then it was ipchains. And actually I'm more familiar with ipchains. I messed with it back in those days.

Steve: As did I.

Leo: Currently, yeah, currently it's iptables. When I heard you say "ipchains," I knew what you meant. I don't think that's a big deal, but there's the correction.

Steve: For what it's worth, if anyone - I thought it was - the only reason, actually the additional reason I wanted to mention it is that what iptables, assuming you are at Linux 2.6 and later...

Leo: And you should be because there's horrible holes in earlier versions.

Steve: Yes. It does allow you the flexibility of accepting traffic on an incoming port and translating the packets so that they essentially are going to, they're mapped to a different listening port. That's cool because, if somebody had, for example, iptables running along with OpenVPN on their own little blue box fanless consumer router, one of the tricks to making sure you're able to reach your OpenVPN server running at home is to have it listening on many different ports because, for example, you might be somewhere which is blocking the default OpenVPN port so you can't get out. But it will allow you certainly to surf to a web server.

So you want OpenVPN also listening on port 80. You're not going to run a web server. You're actually running OpenVPN. But what iptables allows you to do is to have incoming, essentially to listen on port 80 and accept incoming traffic on port 80 and map that over to 1190 something. I think I'm remembering 1191 or 1194, whatever the standard OpenVPN port is. And the idea is you could do this multiple times. So from the outside your IP address is listening to a number of different ports. And that maximizes your likelihood of being able to, wherever you happen to be, get a connection through to your home router. And then that of course allows you to get out to the Internet.

So iptables is what you want. I was just worried that someone might be digging around looking for ipchains and go, oh, shoot, I don't have it. Well, I meant iptables. So thank you, Dan.

Leo: Everybody knows that. You know what, that's how you know you're a Linux

veteran, when you say "ipchains."

Steve: [Laughing]

Leo: Or what was the bootloader before GRUB? God, the one we all used for years. I can't remember. But, see, when you say stuff like that, they go LILO. Yeah. If you say, hey, I use LILO, and we know you mean GRUB, but it just makes you sound like a vet.

Steve: Actually I still have LILO on my - yeah.

Leo: Yeah, you see? Let's see here.

Steve: Because this machine goes running forever.

Leo: [Geezer mode] Yeah. That's right. LILO Loader. [Normal mode] Keenin in Lynnwood, Washington thinks that 50 digits - 50 digits - see, I've been listening to rap music too much. 50 digits should be fine. Give me 50 digits. Steve and Leo, in Episode 403 you mentioned 3.7 X 10^50 when talking about the apparent minimum number - what's wrong with me? - of possible keys available in BitTorrent's new Dropbox-like service. You know what, Steve? That's a really, really big number. To put some perspective on how big a number with 50 digits is, 39 digits of precision in the value of Pi is sufficient to calculate the volume of the known universe to the precision of one atom. Wow.

[en.wikipedia.org/wiki/Pi#Motivations_for_computing_.CF.80]

Steve: Yeah.

Leo: That's obscure knowledge.

Steve: Isn't that cool?

Leo: Yeah.

Steve: That just - that one caught me by surprise. I said, okay, we've got to share this with our listeners. We don't yet know, as I mentioned before, what the protocol for BitTorrent Sync is. And I went looking just to make sure last night they hadn't posted it behind my back and hadn't bothered to tell me. I really think their communications guy was so good when I first talked to him that he will drop me a note when they have a spec ready. And they have said they're going to make that public, and they need to.

What we know is that at least part of - we know that somehow it uses 256-bit keys, but

that some user-controllable part of that is 168 bits. So that's 2^168 possible user controlled-bits. And we don't know what the rest are, where they get the additional bits to pad that out to 256. That's awaiting the spec. But my point was that two - and what I said on the podcast, 2^168 is approximately equal to 3.7 X 10^50 because we can all visualize decimal numbers maybe more easily than bits. 168 ones and zero bits, we know that's a lot. But we also know 3.7 X 10^50 is a number with 50 zeroes.

So what I liked about what Keenin said was that you only need 39 decimal digits in the precision of the value of Pi to be accurate enough, as you said, Leo, as you read, to compute, to calculate the exact volume of the known universe to within the precision of a single atom. So we've got 11 more digits precision than that. And I think that would allow people to hide their user-defined password without worrying about collision. And that's the concern people have, is how do we know we're not going to collide? And it's just there are so many of them, the likelihood is vanishingly small. And we'll wait to see, we really need to see more about the protocol.

Leo: Everybody in the chatroom thinks I'm having a stroke, and that's why I can't talk properly.

Steve: Oh, Leo.

Leo: But I can tell you - and this, actually everybody should know this. And as you and I age, Steve, we should know this, and our loved ones should know this, the three signs of a stroke. If you think somebody is having a stroke, remember STR. You should ask them to smile because, if you have a stroke it'll be like half - it looks funny.

Steve: You're losing control of your facial muscles.

Leo: Right. Ask them to speak a simple sentence.

Steve: Sally sells sea shells by the sea shore.

Leo: And then - that's S T, Smile Talk. And then R, ask them to raise both arms. Okay? So I'm okay. I'm not having a stroke. And then if you want another one, if you want extra help, ask them to stick out their tongue. And then if the tongue is crooked, that's also another indicator.

Steve: Ooh, interesting. How about maybe...

Leo: Okay? Everybody should know these. Pardon me?

Steve: Yeah, let's not have strokes.

Leo: No, let's not. But the reason I mention this is because, like a heart attack, prompt intervention makes a huge difference.

Steve: Yes, it does.

Leo: In this article the doctor said, if I can get a stroke victim to the hospital within an hour of the stroke, I can reverse all damage 100%.

Steve: Wow.

Leo: So there you go. So I just gave you a very important public service.

Steve: And you sound fine, Leo, now.

Leo: And I'm not having a stroke, I don't think. I hope. But I have lost - in the process of doing that, I lost the questions. So let's get them back here. Hold on a second. We were on No. 6 - 7, here we go.

Steve: Yes, we're now on 7.

Leo: David Merillat - I'm going to say it like that. Could be Merillat - in Rochester, New York has some thoughts about carrier-grade NAT and privacy from the RIAA: Steve, I was just listening to today's podcast and heard you deal with the question, "Will CG [or carrier-grade] NAT [network address translation] shield my IP address from the RIAA [the Recording Industry Association of America]?" I think I've got all the acronyms. Oh, no, IP is Internet Protocol. There we go. I think you missed an - well, there were four acronyms in that sentence.

Steve: And you want to prove to us that you have not...

Leo: And I have not had a stroke.

Steve: That's right.

Leo: Holy cow.

Steve: We're with you now.

Leo: Will CG NAT shield my IP address from the RIAA? I think you missed an

important point in your discussion. ISPs already share IP addresses between their customers using DHCP. That's true. Unless you have a static IP, you may get an IP address change at any time. But only one customer can use any given IP address within the DHCP lease period, often 24 hours. ISPs are switching to CG NAT merely because they can't get enough public IP addresses anymore, so they need to share a single IP address among several customers simultaneously. Just as your router does in your home.

So in order to determine where a connection came from down to the level of an individual ISP number, the RIAA will need to record, not just the IP address, but also the source port, which will be chosen by the CG NAT. The ISPs in turn will have to log, not just what public IP address their customer was using, but what ports, as well. If the ISP assigns each port dynamically when a connection is established, this means they would have to log each connection and would actually be able to verify the individual connection that the RIAA was complaining about.

Another CG NAT strategy is to assign each customer a range of ports for an IP address. With this strategy, the ISP doesn't have to log significantly more information than they have to do without CG NAT. In either event, this level of logging is quite possible, and so I agree with you. CG NAT doesn't protect peer-to-peer filesharing users from the RIAA. I detect a total of 27 acronyms in that letter.

Steve: Okay. So he makes some great points. And one thing we encountered was interesting, and that was - I think it was in the original announcement, which I think it was Verizon was the first instance where we encountered this. The second one was with BT, British Telecom, was announcing that a class of their customers were going to be doing this. The AT&T carrier-grade NAT made a comment about only using eight customers per IP. And that sort of set off a little ding for me because what they could do is, much as David said, certainly there is a record-keeping process that ISPs now have to record which of their customers have which IPs at which time. Both internal and external, that is, they're having to - the ISP has to know what their, well, if you don't have NAT, the ISP has to know what public IP they have given to their customers at what time.

Once they have NAT, then they need to go to the extra step of mapping, of keeping track of the mapping between customers' internal private IPs and external public IPs because the only thing the RIAA or the MPAA or anyone else will see is the public IP which no longer is assigned to a single customer. But this notion of port ranges I found very interesting because, as we know, for example, Windows only uses 5,000 ports. Although there are - the port number itself is a 16-bit quantity that runs from 1 to 65535. Windows itself begins assigning so-called "client ports" up above the server port space. The server port space is 1 to 1023. So the first connections that Windows starts issuing to client programs making remote requests is port 1024. And it goes from 1024 up to, what is it, 1024 to 6024. That is, the first 5,000 ports. And then it wraps around. But UNIX systems often just keep on going. Windows never has done that for whatever reason. You don't really need...

Leo: I didn't know that. That's really odd.

Steve: Yeah. You don't really need all those ports. The only time you would ever run out is if something was holding lots of connections open. And there's also, when you shut down a TCP connection, there's a wait period which the TCP/IP stack enforces so that you

don't reuse the same port because you might have old packets drifting around the 'Net which come in, and that would confuse the stack with a new connection between the same endpoints.

So there's like - it's called TIME-WAIT. And if you ever do like a netstat in a command prompt, you'll see typically lots of things that say TIME-WAIT. And that's the stack just saying, okay, we're kind of reserving this for a while until we've given it a few minutes. Then we'll free that up again. And those sort of disappear.

But so the point of all this is that it would be very possible, and it'll be interesting as more information comes to light, for an ISP to map sets of their customers to blocks of external ports so that the RIAA, as David says, now would need to record, not just the IP, but the port number. So the good news is they're going to have to do some scurrying around in order to do that because it's no longer the case that just the IP would disambiguate - I got to use that word during the podcast.

Leo: Congratulations.

Steve: So I'm definitely - I'm stroke-free. And so they would need to use port ranges in order to determine which of a set of customers was using that IP at that time. So anyway, there's still a lot we don't know. Interesting speculation about how this is all going to turn out. But it absolutely is true that just the IP will no longer be enough to identify a person behind carrier-grade NAT. And I think that's good. I like having people have privacy.

Leo: Question 8 is Jim, who wishes to remain anonymous - I'm going to have some more chili - in Toronto wonders whether whitelisting means you can ignore patches:

Steve and Leo, I work for a major Canadian bank which shall also remain anonymous. Our branch workstations run Windows in a very tightly locked-down configuration. Non-admin personnel have extremely limited privileges. Internet access is tightly controlled. USB ports will not mount external storage of any kind. DVD drives are disabled. Anti-virus software is very aggressive, and so on. All the configuration is also continuously monitored so, if someone happens to turn off any of the controls, the monitoring software turns it back on within seconds. Our security department reviews all Microsoft patches and ensures they're all pushed down to the workstations.

I was in a meeting recently, and they were talking about implementing whitelisting. Not only will they whitelist software, but also URLs for Internet Explorer. Now, here's my concern: Someone said that with whitelisting and everything so tightly locked down, well, we won't need to apply patches regularly. We can save them up and apply them, oh, I don't know, once every year or so. I don't know - maybe every decade, just depends.

Steve: I've given away the punch line. Go ahead.

Leo: I don't know about you, but that statement scared the willies out of me. You've mentioned on this podcast numerous exploits that are remotely triggered. I'm not

really in the position to question the approach because my responsibility is solely for the applications that run on the desktop. But there are around 1,000 IT professionals in my building alone. The decision on when to apply patches will be made by a different department, and the corporate culture here is very siloed. It could be career suicide for me to try to tell another department how to do their job, even if I'm right in the long run. Am I being over-paranoid, or does whitelisting really reduce the urgency to apply patches?

Steve: First of all, whitelisting is wonderful.

Leo: What is whitelisting?

Steve: The idea is it's - okay. Blacklisting would be you identify programs you definitely do not want to allow to run. Whitelisting is the reverse. Okay. So first, in blacklisting, the idea would be that the default is to run everything except those programs. In whitelisting, the default is to run nothing except the whitelisted programs. So you reverse the model. The good news is that's great for security because random people cannot bring potentially dangerous or privacy- or security-compromising software into the corporate environment and run it. Their computer will say, eh, that's not whitelisted. You need to go to IT department to get permission to run Quicken to do your accounting on your machine at work, for example. So it really gives tremendous control to IT. And it's great for security. But as to whether that means you need to minimize patching? Oh, my goodness, no. They're completely...

Leo: They're unrelated.

Steve: They're orthogonal to each other. They've got nothing to do with each other.

Leo: But we're secure, Steve. We really are.

Steve: Yeah, we're, hey, we whitelisted Adobe Flash. What could go wrong?

Leo: What could possibly go wrong?

Steve: And of course we need to read PDFs. We whitelisted Acrobat.

Leo: Reader, yeah, yeah.

Steve: What's the problem with that?

Leo: No problem. Java, we might need that, too. Let's whitelist that, too.

Steve: Yeah.

Leo: Hey, just to cheer you up. So anyway, your answer is obvious. I don't know how we get this guy - maybe he's a whistleblower. He could send us...

Steve: Well, maybe these examples are so obvious that he can work this through the system, put it into an anonymous note that he leaves on...

Leo: Well, he does say he's in charge of desktop; right? So he could say, look, that's fine, but you must not whitelist Java, Flash, or Reader. And then they'll say...

Steve: Or they're going to whitelist it, but those you must update because we all know those you must update.

Leo: Right, right.

Steve: And maybe they'll kind of get a clue that, oh, gee, even though we whitelisted those, it doesn't mean that we no longer have to update our software.

Leo: It's magic.

Steve: But I do salute this group for whitelisting. It is a pain in the butt because essentially it locks down everyone's machine so that it could only do, like, only run approved things. But it's great for security. However, it's not a panacea. There's another word now that shows I have not had a stroke.

Leo: Wow, you've disambiguated panacea, and I think there's no stroke in your future.

Steve: Hey, so what were you going to say about making me happy?

Leo: Good news. Twitter has just implemented two-step authentication.

Steve: Yay.

Leo: You can turn it on and go to your account settings in Twitter, require a verification code when I sign in. You'll add a phone, and they'll text you a code.

Steve: Nice, nice.

Leo: I wish they'd support Google Authenticator, but they, that's...

Steve: Wait, wait, wait. You're saying they're not?

Leo: Well, I don't know if they are. They text you a code.

Steve: So, okay, so it's just a...

Leo: Well, I'll sign up, and I'll let you know. But right now it's just we text you a code.

Steve: Well, that's better than nothing.

Leo: Yeah.

Steve: Yeah. I think maybe they're worried that - okay. Giving users an option would be nice. But clearly they're looking at the nature of their demographic and saying, well, we can't ask Granny to do Google Authenticator.

Leo: It's not Granny. It's the 14-year-olds that use Twitter. Granny doesn't use Twitter. Granny's on Facebook. We now found that it's really kids...

Steve: 14 year olds are going to end up with embedded OAuth, like, in their wrists at some point.

Leo: Chip 'em. Chip 'em all.

Steve: Exactly.

Leo: Chip 'em all. Another listener requesting anonymity, also in Canada. This is really an international show here. Steve, I wanted to alert you to a unique solution to poor grades on SSL Labs that has been employed by two Canadian banks. I think we mentioned that TD...

Steve: Many times.

Leo: TD was one of the banks that didn't do so well; right? Block SSL Labs from scanning them. Yay. TD Canada Trust (TD.com), and the Canadian Imperial Bank of Commerce (CIBC.com) are preventing SSL Labs from scanning, period. Wow. That's

one way to fix the problem.

Steve: Yeah.

Leo: Initially, SSL Labs would show an alert, at least for CIBC.com, that indicated something along the lines of "We have been asked not to scan this domain any more." Now, the page simply says, "Unable to test the requested hostname." I think this practice, effectively silencing criticism instead of addressing it, is reprehensible.

On the bright side, their major competitors, Bank of Montreal and the Royal Bank of Canada, both get glowing A's. They have not suppressed SSL Labs scans. Shame on TD and CIBC. Wow.

Steve: Yeah. I just, yeah, this is one annoyance with SSL Labs. Well, and it's not an annoyance with them. They're doing a great service. They do offer a "We will agree not to scan" service. I implemented the same thing on ShieldsUP! when I got a complaint from the U.S. Postal Service that employees apparently within the postal network were using ShieldsUP! and complaining about the fact that there were open ports. And so the U.S Postal Service said we don't want you probing our ports. This is a decade ago. But I still have - I have a small list of domains that I have agreed that I - and what I do is I try to punish them. I say we have been - "Administrators of your network have asked us to please not scan their ports. Take that for what it's worth."

But it's sad that TD.com and CIBC.com would choose to block SSL Labs rather than just fix their servers. They're getting F's because they've got SSL - they're still supporting SSL 2.0, and people are - their hair's on fire. I mean, there are so many problems with SSL 2.0 that it's too easy to play interception games with that. Completely different from certificate abuses that we've been talking about. If you have SSL 2.0, SSL Labs gives you a big red F. And that's what these guys are doing. So they said, oh, don't show that anymore. And to their credit, I mean, I understand the position SSL Labs is in. If you've got somebody saying don't scan us, then you have arguably an obligation not to probe them. Sounds like, though, maybe they've gone, the target domains have gone further and are blocking the source IP of SSL Labs' scan.

Leo: They're not using the system that SSL Labs set up, they're just blocking it.

Steve: Right. They're just saying, eh, stop probing us.

Leo: That's easier. By the way, I'm just checking on Twitter, there's a checkbox, "Require verification when I sign in." And it says "You need to add a phone to your Twitter account." This isn't on by default, but if you use Twitter - there have been so many hacks of Twitter. Even my account I think was hacked once. Not due to any flaw of mine, but actually a Twitter support person had a very poor password, and somebody got...

Steve: It'll be interesting to know because I've got a bunch of iPads that all have Twitter, various types of Twitter clients, persistently logged in. It would be interesting to know

whether it's possible to, like, force a reauthentication when you turn this on, or if it's only for logins going forward. We'll know by next week. And I'm sure our listeners who are curious will be finding out when they're hearing this.

Leo: Moving along to what looks like, sadly...

Steve: No. 10.

Leo: ...our last question from Nick Donnelly, also in London, well, via Saigon. So that's a new country. He wonders about BitTorrent Sync: Love the show. Long-time listener, and much more secure for it. By the way, somebody in our chatroom suggested we should have a second podcast called Security How!. And I really like that idea. It's all about what to do.

Steve: Yeah.

Leo: Wouldn't that be good?

Steve: We could just read - we could read a chapter of the NSA document every week.

Leo: Yeah, every week. Nick says: Love the show, long-time listener - oh, yeah, I read that. But I could always read it again. Love the show. Long-time listener. I listened with interest to the piece on BitTorrent Sync, the vast address space and the probability of brute-forcing a key. I suppose the question is based on current technology. And let's say there are a million keys in use. How long would it take to be able to brute-force a key at random and get hold of somebody's files, say using a single laptop, or a 100-machine server farm? A week? A decade? A million years? I also find it helpful to frame the number of keys relative to the number of grains of sand in the world, or atoms in the universe. Which is it? Hey, that's too much work. Here's to the next 400 shows. Steve?

Steve: So this was sort of a placeholder for me to say what I ended up already saying, so I preempted myself earlier, that we just don't know anything more than everything I shared in our first discussion about BitTorrent Sync. I can't wait to give it a podcast because it is such a cool service. What I do think it is safe to say, though, is that these keys are not brute-forceable. The only way to brute-force a key is to - brute-force means try. You try them all. You start at AAAAA, well, start at A, then go through to Z, then, well, and so forth. So you'd have to be brute-forcing a long key.

And the problem is this is inherently a network transaction. The beauty of the fact that there is no central repository, I think that's what unnerves people. They like the idea of, like, this is my username. Does anybody have that yet? Oh, no. Okay. Then that's my username, and here's my password. Well, the fact is we see constantly these databases being breached. So you don't want that to be anywhere. You don't want there to be a central repository of usernames and passwords.

The beauty of this, and this is why, from everything I've seen, this whole system seems

so well designed, is that you come up with this really long, super-secure token which is hashed down to something, we don't know if it's 160 something or 256 or what it is because we don't yet have the spec. But we're hoping that's still on the way. And then you submit that to a distributed directory, which is going to take time. So you don't want there to be a typo because it's going to have to go out into the Internet, across the 'Net, poke around, do whatever it does, and, like, find the other machines which it believes has this key and then put you in touch with them. And then you try to negotiate an agreement based on a shared secret which is this key, and then you connect. There's no way to do that fast.

And so this is not brute-forceable. That's one of the fundamental aspects of this. And I don't think I explained that well enough because we weren't trying to attack this yet. We need to have a spec in order to really say, okay, if we were the bad guys, what would we do? But all appearances are that there isn't a way to use - to, like, test billions and trillions of these things per second. And even if there were, we're at 10^50. So a billion is nine zeroes. A trillion is 12 zeroes. So, okay, take those off, and now you've still got a lot of seconds remaining from your 50 zeroes. So I think we're okay.

Leo: Oh, you can calculate the volume of the universe.

Steve: Down to a single atom.

Leo: Down to a single atom. And all you need is 50 digits.

Steve: Yeah.

Leo: Of Pi.

Steve: Actually 39.

Leo: 39 digits of Pi.

Steve: So 39 digits of Pi. And we've got 50 digits.

Leo: And we got 50.

Steve: Then, yeah, we're okay. Just choose a good, well, actually, I'm not even sure you choose it. You can submit something.

Leo: No, no, no. You can roll it. But it gives you a random choice each time.

Steve: Yes.

Leo: I don't know how random it is. That's another question. But all of these things, you know, potential flaws, but without the spec you just don't know.

Steve: Yup. So we're going to shelve this for the time being until we get a spec. And then we're going to have a wonderful podcast dissecting BitTorrent Sync.

Leo: I love it. I haven't converted to it, but I love the idea. And I have it installed on all my machines. I just haven't - I use Dropbox, and Dropbox works, and it's kind of done, and I don't really need to do anything about it.

Steve: Yeah. And it fits your security needs. Whereas crypto...

Leo: Yeah, I'm not hiding anything.

Steve: Yeah, exactly, exactly. And next week is Cryptocat. We will analyze the security technology, so maybe - I'm not sure how much of a propellerhead spin this will be. But I know we'll have fun talking about an absolutely bulletproof, secure, multiplatform, cross-browser, very nice, from everything I've seen, well engineered, super-secure chat system, Cryptocat. Crypto.cat.

Leo: Interesting. I can't wait.

Steve: Episode 406.

Leo: Yeah. I mean, I don't need secure chat, but...

Steve: Yeah. Again, lots of people want it. And so I think we're going to have it.

Leo: Yeah. I use, you know, the other issue is that Google has now rolled out Hangouts in a global way. And that's a message system I think a lot of people will be using. What is .cat? Do you know?

Steve: Good question. I don't know what top-level domain that is.

Leo: What country would that be?

Steve: It's kind of cool that it - yeah. I like that their - there's one is cr.ry.crypto.

Leo: Cryp.to, yeah.

Steve: Yes, exactly.

Leo: And "to" is Tonga.

Steve: Yes, exactly. And so I thought that was...

Leo: .Cat is Catalan. Which is a province in Spain, but is also a - it's really intended to highlight the Catalan culture and language. Barcelona is in Catalonia. And it's a different language. It's not Spanish.

Steve: Very cool that [indiscernible] .crypto.

Leo: Yeah. It's not territorial. It applies to the whole Catalan-speaking community.

Steve: And I guess we have Castilian, and that must be a different one, also.

Leo: Castilian is - it's traditional Spanish. And Catalan is really not that much like - it's Spanish-ish. But it's very different. And if you go to Barcelona...

Steve: That's why people are keeping it alive.

Leo: Oh, you're going to get in trouble. There's a big Catalan separatist movement. You don't want to mess with the Catalans. They're great. And believe me, if you ever go to Barcelona, you'll have a deeper under- you should go. You and Jenny. Can you take a week off and go to Barcelona?

Steve: She's my traveling girl, so...

Leo: She will love this. She's been. She must have been.

Steve: She's in New York right now and has been having a ball on Broadway. So...

Leo: Oh, now, see, why didn't you go with her? Steve. You don't like musical theater? You're not a musical theater guy, are you.

Steve: Although I have to say...

Leo: You'd love "Spamalot." You'd love "The Book of Mormon."

Steve: We both enjoyed "The Great Gatsby," speaking of...

Leo: Oh, you liked it.

Steve: Yeah. I mean, I didn't love it. I loved Star Trek. But I liked it. I liked it. I thought it was fun. I think I liked the first half better than the second half. It kind of slowed down.

Leo: Yeah, Sarah Lane fell asleep in the second half. By the way, there is a...

Steve: She's on some strange drugs right now. She's like...

Leo: Maybe that's it.

Steve: She's like - she's like, it was so funny...

Leo: She's like me, probably. She's having a stroke.

Steve: Staring at the screen.

Leo: It's Meow.cat. Look at this. Meow.cat. There's a website.

Steve: Aw, cool.

Leo: That's a name. All right.

Steve: Okay.

Leo: Thank you, sir. This is...

Steve: Goodbye.

Leo: Well, not quite yet. I do want to mention that we do Security Now!...

Steve: Oh, please do.

Leo: ...every Wednesday, 11:00 a.m. Pacific, 2:00 p.m. Eastern time.

Steve: Help me pay my bills.

Leo: 1800 UTC. You should go to GRC.com and buy a copy of SpinRite. Help Steve pay his bills and get a yabba-dabba doo as a thank-you in his office there. We heard a Star Trek sound earlier in the show.

Steve: Ah, yes, that's the hailing whistle. Yes, I had that set to one of my incoming emails.

Leo: [Whistling]

Steve: Yup.

Leo: That's a good idea. I'm going to do that on my phone, have that be incoming email. I like that.

Steve: That's good.

Leo: Yeah. Let's see. You can also get, while you're at GRC.com, 16Kb versions of the show, that's the smallest audio format, and text transcriptions written by an actual human being, Elaine Farris, who does a great job. Hi, Elaine.

Steve: She's actually superhuman, Leo.

Leo: Superhuman.

Steve: Really, yeah.

Leo: Fastest fingers in the West. You also can get all sorts of other free tools like, well, ShieldsUP!, now recommended by the NSA. You should put that right on the front: "As recommended by the NSA."

Steve: I guess that's good. I might probably just stick with this.

Leo: If you want full-quality audio, or even video of the show because Steve is good-looking today in his specs, you can get those at our site, TWiT.tv/sn. And the HoneyMonkeys episode of Security Now!, by the way, is SN-002. I thought it was one, but somebody told me in the chatroom it was two. [SN-001 is "As the Worm Turns."]

Steve: I did, too. And how many times do we regret not numbering from zero? I mean,

that's a mistake you could only make once.

Leo: We can go back and make No. 0 and just pretend.

Steve: Yeah, I don't know what we were thinking, yeah.

Leo: And of course wherever podcasts are you could subscribe to, like iTunes you can get a copy automatically every week when we put it out, which would be very nice. We will see you next week, next Wednesday, Mr. Gibson.

Steve: Yes, you will.

Leo: And with any luck we'll be talking about Cryptocat.

Steve: I hope so.

Leo: If we're not talking about it, something horrible has happened in security and privacy. Thanks, Steve.

Steve: Thanks, Leo.