Listener Feedback #179

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-434.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-434-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. He's got questions; he's got answers. We'll see if we can get through them. Lots of security news, too, including Patch Tuesday notes. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 434, recorded December 11th, 2013: Your questions, Steve's answers, #179.

It's time for Security Now!, the show that protects you, your lived ones, and your privacy online with the man, the myth, the legend, James Tiberius Gibson. No, wait a minute. I'm confusing you with Captain Kirk.

Steve Gibson: James Caffeinated Gibson.

Leo: Steve Caffeinated Gibson. Well, wait, hey, everybody take a drink of caffeinated coffee. Mmm. Hi, Steve. I guess we're getting a little close to the holidays here. Happy holidays to you.

Steve: Yeah, it's been cold. Wow. I mean, like, really cold.

Leo: We've got winter in Irvine.

Steve: Of course you and I are pussies, Leo.

Leo: I know.

Steve: It's like, 50 degrees, oh, my goodness.

Leo: I'm freezing.

Steve: No, it's negative something in [indiscernible] right now.

Leo: You should be in Minnesota right now, where the wind chill is -27.

Steve: So the fates have sort of been good to us because this week did not overload us with stuff to talk about. That is to say, we actually have time for some Q&A. So I'm not sure how much, but we'll certainly give our audience a great probably two hours, if we go the way we have been lately, or nearly that, a hundred minutes or so of podcast, full of interesting stuff and feedback from our customers.

We've got a Patch Tuesday from Microsoft and Adobe. We've got a new release of Firefox. The NSA was found, in another drop of Snowden's never-ending dribble of slides, to be apparently using Google's own cookies to track users. So we'll talk about that. There was a major botnet control center discovered with 2 million passwords, which gave us an additional look into the passwords people are using. FreeBSD, which is the UNIX that I use, has decided to stop trusting hardware random number generators that have recently appeared in chips, for reasons of the NSA. A French Certificate Authority was found to have issued an intermediate certificate that was being abused. Whoops. A bunch of miscellaneous stuff, and Q&A. So, I think, a great podcast today.

Leo: Busy, busy, busy, Mr. Gibson. Okay. If you take a drink, I have to take a drink.

Steve: And I've been putting this together since 5:00 a.m.

Leo: Holy cow.

Steve: I got up at 5:00 to work on the podcast. So six hours of production.

Leo: We are very grateful for the work you do. By the way, you've inspired us, because of the transcriptions you get done from Elaine, we've decided to do transcriptions of many of our other podcasts because it's just a great way to kind of find stuff in the podcast. It helps with searchability. I use it all the time. I'll Google something that I know you mentioned because it will pull up a transcript, and it'll pull me right to that part of the podcast. It's a really great way to do that. So thank you for inspiring that, and we'll have more information on that later today on our Inside TWiT show. Which reminds me that somebody has very kindly offered to, and I know you probably have it in the notes a little later on, but I just want to mention

this...

Steve: Yeah, we should now because I don't say much more than that. But go ahead.

Leo: Bob Noble is launching a project early next year. He wants to - the series has been flagged for translation into, get this, sign language by IndieUnite.com. That's a free service provided by 1 BillionHex, which is a great name for something. There is a concept demo on YouTube, but it's very rough, and it doesn't include the actual sign language yet. So the rigging and the animation of the sign language, they don't use human hands, they use animation. It's time consuming. So he's working on a way to get this done.

Steve: So they would take the audio podcast, and I think he said starting from No. 1, and they felt because it represented such a repository of knowledge that it ought to be available to people who could not hear them.

Leo: He's going to animate each episode from scratch. So instead of using picture-in-picture - but it sounds like a lot of work. Anyway, Bob, thank you. It is a great project. The guy was an IT guy, and he said, "Nine years ago I tried to incorporate some art and acting into my work." And he's always been an advocate of accessibility. So he started by doing talking books for the blind. He got a degree in performing arts with an emphasis on voice acting. Some health problems got him in a wheelchair, unfortunately. And he became the chief of a small, or the chair, that's a good name, the chair of a small charity dedicated to helping indies and organizations better themselves through media and content publishing.

Steve: And he mentioned that, because we have Creative Commons copyright license, that this kind of reuse is permissible.

Leo: Right. We encourage this. Our license - and somebody else asked me if he could rebroadcast our show. I said yes, just look at the bottom of every page on TWiT.tv for our Creative Commons license. Shorthand for it, it's noncommercial, free for noncommercial use. You must give us attribution, and you must reshare anything you do in the same license. Share-alike, it's called. Now, that way somebody can't take this, put sign language on it, charge for it, or put a paywall behind or whatever. And I think that's really, really great.

Anyway, Bob, thank you. We really appreciate it. 1 BillionHex, that's his handle on SoundCloud.com. He also has a CD Baby ShowLink.in/Bob, so he does music, too. Really neat guy. Bob, we'll be in closer touch. But a heads up, look for that coming, and we'll talk about it when it comes out, coming next year.

Steve: And we ought to mention also, we talked about this before we began the podcast, Leo, that I spent the weekend assembling the special holiday episode.

Leo: Oh, baby. I'm excited.

Steve: And for whatever reason, well, first of all, I mean, it was a big deal for me 15 years ago when you and I met for the first time.

Leo: Big deal for me.

Steve: I have it on tape. It was June 17th of 1998.

Leo: Holy cow.

Steve: Fifteen years ago. My hair was dark and...

Leo: My stomach was small.

Steve: Well, so we'll call these our "time capsule episodes" because I took three episodes, the very first three that I had, June 17th of '98. And then a little bit later, actually December of '99, and we were talking about the Click of Death back then with the Iomega Zip and Jazz drives.

Then in December, about this time in '99, your co-host, Kate Botello, discovered ShieldsUP!. I had just put it up on the 'Net, just created it. And she was over there looking for SpinRite, and she said, what's this? And it turns out that I knew her name and saw her hard drive, and she was completely exposed, I mean, which was what happened to everybody in the beginning because everyone's hard drives were out on the Internet, which was what prompted me to create ShieldsUP!. So that episode is her showing it to you because nobody knew about it. And you were like, wait. Our Steve did this? Gibson? And she says yeah.

And then a little bit later, March 3rd of 2000, I was on your Call For Help show and demonstrating it. And we were talking about personal firewalls, which no one really knew about. No one had them. And I made the prediction back then that someday personal computers were going to have them built in because they were so important. It was like, really? You think so? It's like, yeah, I think so. So anyway, fun commercials from back then.

Anyway, it's an hour where I just clipped out, like, the good parts. And some stuff were kind of dumb, but there's like a newscaster predicting that someday - that Nokia made a prediction that somebody handheld phones may be connected to the Internet. It's like, oh, my god. So I think everyone's going to get a kick out of it.

Leo: I can't wait. That is a lot of fun. All right. Well, we have lots to do today. It's a busy day. We've got a Q&A, first time in a month. And of course I should mention that show will air on Christmas Day, December 25th, in place of our regular show, at the usual time.

Steve: And like last year's, it's really one, you know, you could listen to it; but, boy, the visuals are something you're not going to want to miss. So since people will be home, maybe they can break their regular commute cycle and watch this one.

Leo: Download this one on video, yeah, yeah.

Steve: Yeah. This one you're going to want to see.

Leo: 11:00 a.m. Pacific, as always, 2:00 p.m. Eastern time, 19:00 UTC, Christmas Day. Get up in the morning, open your presents, and watch Security Now!. Believe it or not, Steve, there are people that will do that.

Steve: So is that what you're going to do? You're going to air these specials and best-ofs and things at the same timeslots they would normally be live?

Leo: Yeah.

Steve: Okay, cool.

Leo: I think that's the plan. You know, I haven't asked, but I think that's what we will - normally that's what we do, yeah. So if you're in the habit - and a lot of people don't observe Christmas. So if you're in the habit, and it's a Wednesday, it's a normal day for you, and you tune in, you will get Security Now!, just a very special Security Now!. And many of our other shows are doing best-ofs. And of course don't forget the following week is New Year's Day. We're doing our 24 Hours of New Years. I've been trying to convince the staff for years to do this. They finally acceded. Now that I'm an old man, I don't know how I'm going to do it. But starting 4:00 a.m. Pacific, 7:00 a.m. Eastern on New Year's Eve, that is 20, no, I'm sorry, that's 12 Noon New Year's Eve Day UTC, and then going to 4:00 a.m. New Year's Day. So 4:00 a.m. New Year's Eve...

Steve: And you and I will be...

Leo: And you're going to come up; right?

Steve: You and I will be together, yep.

Leo: So we've got two - we've got a normal Security Now! next week, and then it gets weird. Christmas Day and New Year's Day. It's going to be a lot of fun. And I should still be alert by that time. I won't be passed out yet. Or actually, wait a minute. New Year's, no, you're New Year's Day. So I don't know what we're going to do for Security Now! on that Wednesday. I will have gone home. I'll be fast asleep.

Steve: Yeah, we're doing it on Tuesday. We talked about this before.

Leo: So we'll do it ahead of time, and then - yeah.

Steve: Right.

Leo: Okay. All right. The best part of this, though, is you're going to get 26 five, four, three, two, one countdowns, Happy 2014, because we're going to start, and I think it's Papua, New Guinea, whatever the Western, or is it Easternmost, whatever the...

Steve: You said there were some half-hour time zones, too.

Leo: There's some 15-minute and half-hour time zones. I'm told there's 27 time zones in the world. And we're going to try to hit each one and count down each one. That's a lot of champagne, 27 glasses of champagne. Wow. You know, I've got to...

Steve: I'll be bringing...

Leo: Go ahead.

Steve: I was going to say, I'll be bringing my coffee-fixing stuff up, and we'll have you tasting...

Leo: We'll have coffee.

Steve: We'll see, yeah, we'll...

Leo: I think a lot of coffee, yeah.

Steve: I think so.

Leo: And I think I may not drink champagne. I might drink Martinelli's Sparkling Apple Cider. It just looks like champagne.

Steve: Yeah, because, you know, you need stability to be on that ball, Leo. The ball, you need...

Leo: Oh, I ain't doing 24 hours on a ball. I'm not nuts. All right, Steve. What's the

news across the nation?

Steve: Well, we are here on our second, we just passed the Second Tuesday of the Month. It's funny because when we announced the podcast will be moving to Tuesday, I got a bunch of tweets from people saying, wait a minute, does that mean you won't be able to do Patch Tuesday anymore? Because we're taking advantage of the fact that we're one day later than Patch Tuesday sometimes. But Microsoft does release this information in some mailings that I subscribe to, so I think we should still be okay with our - we'll be a little more current, in fact, with Patch Tuesday. It won't be yesterday, it'll be today that those patches are available.

We have some important things, both from Microsoft and Adobe, because they're zero-day fixes, meaning that they were discovered because bad guys were found already doing them in the wild before the vulnerability was known. So these are, as always, you want to stay current. There are 11 patches fixing 24 vulnerabilities. So that's kind of medium size, relative to what we see on these monthly updates from Microsoft. Twelve of them, 12 of the 24, so half of them, were remote code execution vulnerabilities. Those are never good.

You have eight elevation-of-privilege vulnerabilities, which is where, if you run an unprivileged account for security, which of course is always what people should do, this is a way of essentially getting admin privileges when the sandbox essentially that you're running in with deliberately limited privileges doesn't allow those, so those are not good. Eight of these are being fixed. And they call them denial-of-service vulnerabilities in two cases. They're not what we think of like in terms of a bandwidth flood. They're basically a crash. Something is able to crash some software, thus denying you the service of that thing it crashed, whatever it is. So two of those got fixed.

An information disclosure vulnerability, that's one that people are considering worrisome because it involves the cloud services and the way Microsoft's cloud services function, which would allow account information to escape. So you want to pay attention to that one, too. I think that was 104. And then a security bypass vulnerability.

So what Microsoft is now doing is ranking these, as we've talked about before, in the order in which they should be done if for some reason you can't do them all. And they're sort of sequential: 96, 97, 98, 99, and 105. They skip a bit. Those are the, like, if you only can do some, do those to immediately prevent exploitation by attackers. And then 96 of that, the very first one, is the critical zero-day vulnerability in Windows and Office; 97 and 99 of that group fix a dangerous scripting problem in Windows. And Microsoft has said, of those three, of 96, 97, and 99, they're expecting active exploits immediately, and in some case that's already happened. So those need to get done. And then the lesser important ones - 100, 101, 102, 104, and 106 - should be ASAP, like as soon as possible. And then 103 is like, eh, at your earliest convenience, says Microsoft. So anyway, do them all. End users certainly should. And I guess admins who don't have a choice, do the important ones.

Leo: Right.

Steve: Adobe's updated Flash and Shockwave, fixing two security holes, including one that is in the wild now, being actively used in attacks, malicious Shockwave Flash .swf files are being attached to Microsoft Word documents and emailed to people. And the

curious who open the Word doc get themselves exploited immediately.

Leo: Is that using a macro? How are they doing that?

Steve: It's just an attachment. And I guess the act of opening it when it's attached to the DOC file will cause it to get executed.

Leo: Wow.

Steve: So that's a question. It may well be a macro that is saying, oh, run this. So the DOC says run the attachment, and then the attachment says, oh, thank you, and takes over your computer.

Leo: The reason I ask is because Word now won't run a macro automatically. It'll have to - it'll say, do you want to run this macro? I wonder how they're doing that. It's interesting.

Steve: Yeah, they probably get around that.

Leo: Yeah, maybe they got around that. Yeah, there you go.

Steve: You know, it's like the address space layout randomization, that malware just says, okay, yeah, well, you have that on, but we're going to get around that.

[Talking simultaneously]

Steve: We have a new Firefox. That happened a couple days ago. I announced it to my Twitter followers. Not a big deal. The one really nice thing - this is v26, so anybody who wants it, who just, like, leaves Firefox running all the time, as I do, go to About > Help, or Help > About, rather, in your Firefox menu, and up comes the About box, and you'll see that it immediately kicks it into downloading the update. And then you need to restart Firefox. Which is fine because it, like, brings all your tabs back, and you'll have the latest and greatest.

What they did was all Java plugins are now click-to-play. So I don't know what took them so long. But that's now in v26. That's a big change there. So that, if any Java is part of a web page, it will not run by itself. Which, yay, that's absolutely what we want. You have to - it'll show it to you disabled, and you have to explicitly click on it in order to run it. Which is a reasonable tradeoff. We know that there are sites that depend upon Java, that won't run without it. But running without intervention is a massive, as we've often talked, security vulnerability. So now you just have to click on it in order to run it. So it's like, hey, that seems reasonable.

They've also updated their password manager so it now supports script-generated, that is, their built-in password manager in Firefox, so it now supports script-generated password fields which it just wasn't aware of before. They made some changes to the way Update works with Windows, so that the user running the update no longer requires

write permissions as long as they have, like, a Mozilla maintenance service which runs in the background. Because it's a service, it's able to have system-wide rights. And so they've arranged for the non-privileged code, which sees the update, to communicate with the service and get the service to do the update so a non-privileged user is able to get their Firefox updated nicely.

They added also H.264, which of course is now the new standard video codec for Linux, which didn't have it in Firefox before. And then there's a bunch of developer improvements and miscellaneous fixes. So anyway, it's a good thing to get, and it's easy. Just restart Firefox after you go to Help > About.

Okay. So disturbing news from the next Snowden dribble. And I have to just take my hat off to the strategy because, first of all, it is, when you think about it, it is phenomenal how much Edward got before he left, the fact that even now we're still getting new revelations from the data that he collected. And it's been so much more effective than if he'd just dumped it out on the world and said, "Here," because we would have been overwhelmed. It would have been, oh, my god, and the news would have been significant for a week, and then it would have been forgotten. So this is the way to do it.

What we've learned, and the Washington Post covered it a couple days ago, is that in a slide that was recently released, that it made reference to the NSA and GCHQ, the equivalent agency in the U.K., using cookies, and specifically the Google PREF ID cookie, and this slide says to enable remote exploitation. And it is not clear how that exactly works because, again, a lot of interpretation needs to be done of these slides. But the Washington Post wrote, they said: "The agency's internal presentation slides, provided by former NSA contractor Edward Snowden," which is the standard byline now for these, "show that when companies follow consumers on the Internet to better serve them advertising, the technique opens the door for similar tracking by the government," essentially piggybacking on the tracking that other trackers are doing.

"The slides also suggest that the agency is using these tracking techniques to help identify targets for offensive hacking operations." And it talks about enabling remote exploitation. So my take on this, as I started saying, is, well, okay. We have to assume, because all the evidence now demonstrates that what can be done is being done, the NSA are full of really smart people, and so they're as able to look at traffic on the Internet as hackers can. And they can look at this, I mean, and their bosses are saying, "We want you to track everything. You know, everything. So do that."

So they're seeing users' browsers sending out cookies with every query that the browser is making. And they're saying, well, okay, why can't we track these? And the answer is they can. If they're sitting here looking at a big pipe, and they're able to obtain nonencrypted, which is to say non-SSL, standard HTTP queries, then all of those queries are containing cookies wherever they're going.

I mean, I'm sure you're aware, Leo, how pervasive Google's Analytics are on websites? I mean, Google Analytics is, like, everywhere because a phenomenal number of sites use it. Well, that means that there is Google script running pervasively across the Internet, and that script is making queries to Google which are sending the Google Analytics cookies back to Google. And anybody sitting on the Internet looking at all the unencrypted traffic can absolutely track users as easily as Google can. They're getting all the information that Google is, and potentially all the information that all the tracking advertisers are getting. I mean, so why wouldn't they be doing this? It makes perfect sense that that's what someone would be doing in order to aggregate this.

This sort of makes the NSA like this super cookie tracking organization because it's cross

organization whenever the connections are unencrypted, which unfortunately is still a lot of the time. We've got a question later on from someone who installed Calomel and was a little shocked by how many sites had no encryption whatsoever. He just assumed that everybody was doing it. But many sites are still not doing it. So it's like, yeah. Again, we ought to assume what can be done, the NSA will be doing.

And the fact is there's so much use of nonencrypted traffic, where encryption is still only used during password negotiation, when you negotiate the password. And the problem, of course, as we know from our coverage of Firesheep earlier, and if any listener doesn't know about Firesheep, go find that podcast because it is still happening [SN-272]. Anytime you use a secure connection to log in, but then your browsing and your movement around the site subsequently is not SSL, that means that the way you are maintained, your login is maintained, is a cookie is going back and forth to the server in the clear.

And what Firesheep showed was how easy it is to obtain those cookies in any open WiFi setting because anyone on the WiFi, it's like you're on a hub on a network. You're on a shared network. And you can see all of the cookies that everyone is using. That is their logged-on session. And so it is trivial for someone to just start making queries to the same server and giving it that cookie. They're logged on, too, as the person whose cookie they stole. And this is still very pervasive. So, I mean, it really does say that we need to get HTTPS all the time, everywhere. And some major websites, as we'll find out later in this podcast, are still not doing that.

Leo: By the way, Snowden one of 10 people nominated for Time magazine Man of the Year. Didn't win it. Pope won it. There was some pretty tough competition.

Steve: That's pretty stiff competition, yeah. He seems to be a neat Pope, too, so...

Leo: He's a good Pope. I probably would have preferred Edward Snowden, but it's not up to me.

Steve: That's a little dicey because...

Leo: Well, remember, Man of the Year is not a good man, necessarily. Hitler was Man of the Year.

Steve: Oh, okay.

Leo: It's just the most important person, newsworthy wise, that year. Right? So it doesn't have to be good.

Steve: In that case I agree, Leo. I think, with that criteria, Snowden really does deserve it.

Leo: You tell me. Here's the candidates: Bashar Assad, President of Syria; Jeff Bezos, Amazon; Ted Cruz, Senator from Texas; Miley Cyrus. I don't know how she got on that list. For good or evil, I don't think she made a difference. Pope Francis; President Obama; Hassan Rouhani, President of Iran. You see, it doesn't have to be somebody good, just somebody who made a big difference; right?

Steve: Yeah, yeah.

Leo: Kathleen Sebelius, Secretary of HHS.

Steve: Whoops.

Leo: Not Man of the Year, I don't think. But Edward Snowden, NSA leaker; and Edith Windsor, the gay rights activist whose Supreme Court case cleared the way for gay marriage. So those are all very important. But, boy...

Steve: Wow, that's some stiff...

Leo: Edward Snowden's, to me, on the top couple or three.

Steve: Yeah, because of the nature of the way he did this.

Leo: He's changed the world. This changed the world.

Steve: Yes. Yes. I didn't cover, because it just didn't seem, I don't know, quite relevant enough, but Silicon Valley is really getting together now and saying we've got to put together a coalition to fight this, to fight this pervasive surveillance. And they've put together, I guess, a letter and sign-up sheet and so forth in order to begin to act. But the NSA was critically hurt because it turns out that what they were doing, as we know, was way more than people feel they really had a license to do.

Leo: Yeah. By the way, even Time magazine calls it Person of the Year. I'm going back in time when I call it Man of the Year. I don't know when they changed it to Person of the Year, but it's Person of the Year. Or Human of the Year.

Steve: Wasn't there going to be a machine? Was it Watson? I thought there was...

Leo: For a while I think Watson was the - what is the history of Time magazine's...

Steve: Entity of the Year, Leo.

Leo: And why Time magazine gets to decide this, I don't - just because they say they do, I guess. Let me see, the Person of the Year.

Steve: Well, yeah, and Forbes has their 500 and their 100 and that.

Leo: It started in 1927, Charles Lindbergh. Last year it was the President. The Protestor was, because of the Arab Spring and Occupy Movement, in 2011. Mark Zuckerberg in 2010. That makes no sense. Ben Bernanke of the Fed, 2009.

Steve: Although remember Facebook was a big lot of noise back then.

Leo: Yeah, yeah. Good Samaritans in 2005: Bono, Bill Gates, and Melinda Gates, kind of a triumvirate. The American Soldier, 2003. You know, in 2002, Whistleblowers were the Persons of the Year. And I think that's what Snowden is; right?

Steve: Oh, my goodness, of course, yes.

Leo: So let's just say he won in 2002 for - collectively. The Pope was picked, Pope John Paul II, in 1994. So it's not the first...

Steve: Oh, no kidding. But a different Pope.

Leo: Different Popes. Francis was just a twinkle in the eye of the College of Cardinals. I'm sorry. I'm going straight to hell. Continue.

Steve: So a security research group, Spider Labs, discovered a massive database of usernames and passwords. And so this is different than the typical vulnerabilities or disclosures that we've seen before. Typically what we report on is a given website has lost control of their database. And so we're looking at the passwords of that website. In this case, what was found was the repository from a botnet. So this is - it's called the Pony, P-o-n-y, the Pony Botnet because it uses - the botnet controller is identified as the Pony controller. And I guess the controller itself is now open source, or the source has been published. And it was by looking at the source that they were able to find this one controller of one botnet.

And so there are bots installed on people's machines, that they're clearly not aware of, that are collecting - the bots are spying on them, collecting account information, their credentials, as they log into various accounts. And so 1,580,000 website login credentials were found; 320,000 email account login credentials; 41,000 FTP account credentials; 3,000 RDP, remote desktop, Windows remote desktop credentials. And that's bad because that means that anybody who's got that can log into your remote desktop, which is full access - typically those are servers - full access to your machine then. Three thousand secure shell account credentials stolen.

So first of all, you've got nearly 2 million login credentials found. But then, as is always interesting to security researchers, they look at what these are. And no surprises. The password 123456, ever the most popular password, was used in 15,820 cases. In fact, I've got a little table there from the breakdown. Second most popular, little bit longer. People thought, well, I'm going to make this password longer. So they added 789 to the end of 123456. Then third most were the little more lazy people. They only went as far as 1234, and then they said, ah, you know, who cares about the 56. That's obvious, so we'll just leave that off. And then the word "password" is No. 4 on the hit parade.

Leo: It always is, rhymes with "assword."

Steve: Exactly.

Leo: Yup.

Steve: And then 12345, that made No. 5 on the list; 12345678, they didn't go as far as 9, but they went a little further than 6, and so that was No. 6. I mean, and so forth. I mean, this is - and I don't know, 1,224 sites accept the numeral "1" as a password. So this is just...

Leo: Wait a minute. One digit?

Steve: One digit.

Leo: Just one?

Steve: One, the numeral "1," and whatever it is says, okay, that's your password. That's fine. That must be email, dumb email servers. I just, you know, wow. So anyway, interesting list. It generated enough news that KTLA TV had me on.

Leo: I heard about that, yeah.

Steve: Last Thursday, I think it was. And Rick Romero is their consumer guy. And his mission, which is why he brings me back and reminds people about Password Haystacks at GRC, is just trying to get people to stop doing 123456, just for their own good. Please do anything but that. And so, which is why he likes the Haystacks idea because I'm just saying, just do something more. Add a bunch of something to it. Do something. Doesn't have to be hard to remember. Just get off this list because that's better than nothing. And it's easy to do.

Leo: You've seen the HaveIBeenPwned website; right? It searches that database. So you can enter in your email address.

Steve: Yes. Nice.

Leo: And the one that I - when I enter my email address, it says Adobe because, of all of the accounts that were hacked, that's the one that - and I knew about that. So, yeah. So if you want to know, it's HaveIBeenPwned.com, HaveIBeenPwned. And, now, I don't know these guys, and it's Troy Hunt. I don't know if he is saving the passwords or whatever, but I...

Steve: Oh, Troy's a good guy. I would trust Troy completely.

Leo: Yeah. It's fun to look. And you should never use the same password everywhere. What is the - the bitcoin's gone down, hasn't it.

Steve: I think it was at nine something last time I looked. But it seems to have stabilized. I mean, it didn't, like, crash all the way back down. It dropped down into the eights or sevens. And so it seems to be sort of holding.

Leo: Good. Good, good, good. As a holder of 50 bitcoins, you probably pay a little bit of attention to this.

Steve: Well, I do love that ZeroBlock app on iOS. Every time - I sort of see it on my list, or in my little screen of apps every so often. I go, oh, where are we now, and I tap it and go, oh, okay. Just sort of, I mean, I'm never going to sell. I'm just going to ride this wherever it goes. I mean, I didn't do anything to earn them. They just sort of appeared magically, back in the days when you could actually have a PC mint for you. Now, I guess, I heard some news that China's getting very interested in bitcoin because they like the idea of a non-state-controlled currency in transactions. But they're, like, setting up warehouses with mining machines. It's like, oh, there goes the neighborhood. So...

Leo: Ugh.

Steve: Yeah. It's getting to be - and apparently there's - I did see someone, there's an organization, a mining pool group that said...

Leo: I saw that.

Steve: ...they owe their existence to us, to Security Now! and...

Leo: Oh, I didn't see that. Really.

Steve: ...talking about, yeah, to talking about bitcoin on the show. That's what created it. So I thought that was cool.

Leo: By the way, somebody in the chatroom's saying you can use OpenVPN on iOS.

Steve: Yes, I thought that was the case. I thought they were now...

Leo: They have both, so Android and iPhone. Sorry.

Steve: So my UNIX is FreeBSD. Brett Glass told me about it years ago. You remember Brett probably from the old days. He's still around. And he said that's the one you want. That's the one I went with, and I've never been unhappy with it. It's what runs our news server, and I use BIND for my DNS server back at the server farm at GRC. And that's my choice. So Ars Technica carried a story that indicated that the next release of FreeBSD, v10, they are backing off of their endorsement and use of hardware random number generators which have recently been incorporated into chips. Which, I mean, it's sad, but I completely understand it. We know that the NSA has attempted to influence chip design. We know, we really suspect, we have to suspect that the NSA influenced an organization as important and significant as RSA to choose the worst, I mean, the known defective pseudorandom number generator as the default for all of their cryptographic library packages. We have to assume that's why they would have done that.

And so here Intel comes out with a much ballyhooed hardware quantum-level uncertainty, I mean, beautiful pseudorandom number, actually true physical random number generator - I'm so used to saying "pseudo." This is not pseudo, it's truly random, and now we can't use it. I mean, it's so sad. But it's true. We can't. So actually what they're doing is they're assuming that it can't be trusted. So they're falling back to what they were doing before, which is to use Yarrow.

Yarrow was designed by our buddies Bruce Schneier, John Kelsey, and Niels Ferguson at Counterpane Labs, Bruce's group. It's a very beautifully designed pseudorandom, software-based so therefore necessarily pseudorandom, number generator, but it works on the concept of a pool of randomness, which is the standard way you do really good random number generators now. I looked into it carefully when I was doing - I needed a really good pseudorandom number generator for the Off The Grid project that I worked on where I used Latin Squares in order to generate web name-based, domain name-based passwords. And I didn't use it because I needed even more entropy. It uses by default a pair of SHA-1 hash contexts. An SHA-1 context is 160 bits. So essentially it sort of bounces back and forth. It uses one while it's building the entropy in the other one. And once the entropy is used up in the one it's using, it sort of ping-pongs. It switches over to the one which has been building up entropy and then begins rebuilding entropy in the other one.

But the problem is you're actually, I mean, that's really good randomness, but I actually needed more. Remember that I had what I called a UHE PRNG, Ultra-High Entropy Pseudorandom Number Generator, that I developed. It uses 1536 bits. And the reason was there were, like, that many Latin Squares. And so if I used any lower entropy random number generator, then I couldn't get to all of the Latin Squares that were possible. So I had to do that. But that's way beyond most systems' needs. They just need, you know, give me a chunk of entropy.

So what FreeBSD is doing is they used to be using a purely software PRNG. Then they switched to a purely hardware PRNG, or sorry, RNG, not pseudo, really, really fabulously random, true random. But now they're having to back off of that. So the hardware

randomness and other sources, since they can't purely trust the hardware random number generator, that will all then feed into Yarrow, which is good because all of these things are, like, hungry for entropy. When we talk about entropy being used up, the good pseudorandom number generators sort of keep track of how much randomness has been taken from the pool. And at some point they decide it goes below a threshold, that's when they switch over to one which has been collecting entropy while the other one's been dispensing it. And so they ping-pong between.

So anyway, this is probably a trend we're going to see. These guys are right. We cannot simply trust the hardware because it is subject to really subtle manipulation. And subtle manipulation is, as we know, all you need. Any deviation from randomness is not good for crypto.

Leo: Somebody in the chatroom is saying you could build your own hardware true random number generator with a few chips and a reverse-bias transistor. That would be a great project. Is it true that it would be effective?

Steve: Yeah. You can use, for example, a reverse-bias diode. And what happens is electrons migrate absolutely randomly across the reverse-biased PN junction of a diode. And then you amplify that and count them, and it generates entropy. Sometimes the problem is it's not generating it rapidly enough. And so that's one - what Intel did was they actually used cross-coupled inverters which could never be in a stable state because they would just - entropically it couldn't be stable. I guess it was three in a row. And it was just, like, screamingly fast. And then you sampled that and got really good entropy.

I would, I mean, I tend to think Intel's probably hasn't been warped, but we can't know. But it is the case that you could absolutely build - there's lots of different approaches now to build that. And there's something called the "entropy key" I've been trying to buy from an outfit in the U.K. now for quite some time, just for some reason it's difficult to get, and they've stopped making it available. But it was just a little USB plugin dongle. And in fact even Yubico has in their - and here, I happen to have it in front of me. This thing is their - I can't remember what they call it. It's one of their gizmos. It's a USB device. It's got crypto in it. But it also has a hardware random number generator built in.

Leo: Is it, like, true random number?

Steve: True random number, yeah. True, based, in hardware. And of course I would trust Stina till the end of the world, not to be...

Leo: So I could plug that in and run some command, and it would spit out a random number. It probably has a program on it; right?

Steve: Yes.

Leo: That's something special, those. That looks like a USB key. I have a bunch of YubiKey stuff.

Steve: Yeah. If we go to Yubico.com, you'll see it's an acronym, HMS or something or other.

Leo: All these Yubico things. I have to figure out which one - this one you press a button and it does something, so that's the normal Yubico. I don't know. I'll have to figure - and then I have one with NFC in it that's really cool, that I just tap stuff to.

Steve: Yup, yup. And in fact Google has announced that they have their U2F, I think, Universal Two-Factor project. So, yeah, this news was a couple weeks ago that Google and Yubico have, like, officially teamed up. Google's using it in-house now. And it uses the NFC-equipped YubiKey to do two-party, two-factor authentication. Of course there's still username and password, so it's not - it doesn't have the same goals that I have with SQRL, which is to completely replace username and password. You'll probably see a gizmo there, if you find it.

Leo: I'm looking through all the different products.

Steve: It's their hardware something for their server.

Leo: I didn't realize they have a dedicated LastPass YubiKey. That's a cool thing. That is neat.

Steve: Yeah, they're doing great. It's really neat this came up...

Leo: They're making relationships with a lot of different people, which is really, really cool.

Steve: Yeah. Now that Stina is over here in Silicon Valley, she's got access to people.

Leo: Yeah. The NSA can't hack those. That's the good news.

Steve: I would really trust...

Leo: Well, if you built your own with a reverse-bias diode, then you know it's okay, unless the NSA got to the diode manufacturer. I don't think that's going to be...

Steve: No, actually they...

Leo: They haven't yet gotten to physics. God is not on the NSA's side, thank goodness. I just emailed Iyaz and Father Robert Ballecer and said let's do this on Know How as a project.

Steve: That'd be a great idea. I'm sure if you Google "build my own hardware random number generator," you'll find a bunch of stuff there. You'd probably - you'd do a simple little reverse-bias junction. Sometimes a tunnel diode is used. And then you'd, like, run that through an Arduino. So it would generate noise. You have to do post-processing. You need to whiten it and balance it and then, like, run your own routine to verify that it's producing good stuff so that it can, like, shut it down if it's not. But it can definitely be done. I'm sure people have done that.

Leo: And this is the - you were talking about the Yubi, it's called the YubiHSM.

Steve: That's the one.

Leo: Hardware Security Module.

Steve: Yeah. And the idea being that you would plug this into a server, and it's able to then - it cannot be - it's like Apple's Secure Store in iOS. It's a deliberately read-only, you can't write to it sort of thing. And all you can do is ask it questions. And so it creates a secure boundary such that you can keep your secrets there, and no compromise of the server is able to compromise it because it's just - it's standing outside with a clearly defined serial interface through USB to the rest of the device. So you can store secrets and things in it.

Leo: Five hundred bucks, so...

Steve: Yeah. Yeah. Okay. So this was a bit of spin from the French government. The news hit on Saturday, on December 7th. Google's online security blog said: "Late on December 3rd," so this was they were blogging after four days - "we became aware of unauthorized digital certificates for several Google domains." Now, remember, this is one of the very cool things that Chrome is doing is Chrome knows what certificates are authentic. And so the instant you try to use Chrome on a site with a certificate that says it's from Google and it should be trusted, Chrome will say, uh, wait, where did you get this certificate, we never produced that, and immediately phones home and raises flags. So this is very cool feedback that Google has built into Chrome.

So Google says: "We investigated immediately and found the certificate was issued by an intermediate certificate authority linked back to ANSSI, a French certificate authority. Intermediate CA certificates carry the full authority of the CA, so anyone who has one can use it to create a certificate for any website they wish to impersonate." So that's - we'll talk about why this was done versus the normal way in a second.

And Google continues, saying: "In response, we updated Chrome's certificate revocation metadata immediately to block that intermediate CA, and then alerted ANSSI and other browser vendors. Our actions addressed the immediate problem for our users. ANSSI has found that the intermediate CA certificate" - now, that's what they issued. So they issued this with full recertification authority because they could constrain any certificate they issue, but they said we're issuing an intermediate certificate that can itself issue its own certificates. So that was deliberate.

They said it "was used in a commercial device, on a private network, to inspect encrypted

traffic with the knowledge of the users on that network. This was a violation," they're now saying, "of their procedures, and they have asked for the certificate in question to be revoked by browsers." Uh-huh. "We updated Chrome's revocation metadata," says Google, "again to implement this. This incident represents a serious breach and demonstrates why Certificate Transparency, which we developed in 2011 and have been advocating for since, is so critical. Since our priority is the security and privacy of our users, we are carefully considering what additional actions may be necessary."

Now, separately, same day, ANSSI posted on their site - this is the Agence nationale de la scurit des systmes d'information, so that's where ANSSI comes from. They said: "As a result of a human error," which we'll discuss in a minute, "which was made during a process aimed at strengthening the overall IT security of the French Ministry of Finance, digital certificates related to third-party domains which do not belong to the French administration have been signed by a certification authority of the DG Trsor (Treasury) which is attached to the IGC/A. The mistake has had no consequences on the overall network security, either for the French administration or the general public." Well, the general public might beg to differ. But they continue: "The aforementioned branch of the IGC/A has been revoked preventively." Yeah, because it doesn't work anymore because no browsers will honor it.

Leo: Preventative revocation.

Steve: Yes "The reinforcement of the whole IGC/A process is currently under supervision to make sure no incident of this kind will ever happen again." Okay. My comment: This could not have been human error. Human error is when it doesn't work the way you want it to. Deliberate function is when it does. This had to be deliberate. What we know is...

Leo: Ah, whoa.

Steve: It had to be deliberate. There are two ways that an appliance can function, and we've talked about them. The way you do an appliance which wants to have visibility into the network traffic is you have only two choices. You install your own certificate into every browser which will trust that appliance. And of course that's burdensome. That means everyone who wants to use the network will have to install a certificate to trust the certificate that the appliance has because it has to mint certificates on the fly which it signs, so you need every browser to trust the certificates it mints.

Well, gee, that's awkward. Wouldn't it be nicer if instead we minted certificates that a trusted intermediate authority is issuing, and it's trusted because a true root authority that all browsers already trust are trusting? And the answer is, well, of course that would be much nicer. Then we don't have to install anything in everyone's browser. So somebody deliberately configured this appliance to work that way. This cannot possibly have been a mistake. If, I mean, whoever did this knew exactly what they were doing, that they had somehow acquired an intermediate certificate with full certificate signing authority, which gave it completely unrestrained ability to create certificates. So now everyone behind that appliance was having their traffic inspected, full spying on all their traffic with no warnings from their browser and no need to install a certificate. Any corporate environment that is behind such an appliance must have, must trust the certificates that the appliance is minting. So it's got to accept that in its browser. This avoided that.

So this wasn't a mistake. This wasn't inadvertent. They got caught is essentially what happened because somebody ran Chrome inside that network. And running Chrome is all you have to do to shut that down instantly. Chrome will scream home to Google that somebody has given it a Google cert that did not come from Google, and that's end of the game.

Leo: Wow.

Steve: Yeah. Very cool. Yeah, but not a mistake. They got caught.

Leo: Fascinating. And they didn't admit it, which really...

Steve: No, absolutely.

Leo: I think that breaches trust, and I think...

Steve: Pure CYA.

Leo: Yeah. And it's a shame. I mean, when you get busted, just admit it.

Steve: I don't know how they could, though. I mean, again - and of course this is why none of us believed the heartfelt "We're not spying" after the first round of NSA revelations. It's like, well, you have to say you're not. You don't have a choice.

Leo: And of course you assume that no one is going to be sophisticated enough to call you on it except Steve Gibson and a few others.

Steve: You read this, it's like, oh, okay, that sounds fine.

Leo: Yeah, the normal media is just going to go, oh, yeah, wasn't their fault, must have been a...

Steve: Somebody pushed the wrong button somewhere.

Leo: Just pushed a wrong button.

Steve: Yeah. So Forbes had an interesting story, and I don't know what to make of this. But Apple just rolled out their new iBeacon technology.

Leo: Oh, we've been talking a lot about this, yeah. I wanted to get your security take on this, for sure.

Steve: Yeah, makes me very nervous. I mean, just Bluetooth is - I don't like Bluetooth, and I don't like NFC because radio is scary. I mean, radio is very powerful when you use it right. But it's like, I mean, and Bluetooth I love. We've talked about Bluetooth. The technology is solid. But it just - you have to be very careful.

Leo: And this uses Bluetooth LE, which is - to me it's interesting because it doesn't require pairing. One of the security features of Bluetooth was you have to have explicit pairing between two objects. You can exchange passcodes. You really have some form of security. But Bluetooth LE doesn't work that way. If you have a Bluetooth LE device in an application, you launch the application, it sees the device. And that's how iBeacon works.

Steve: Yes, exactly. So the idea is like all over the store you've got little beacons. And so as you walk to them with your phone, your phone is able to receive information from them. It's like, oh, look. And so what was interesting was that Forbes noted that in the Objective-C interface description for the iBeacon framework, it says this technology, this framework allows you to scan for Bluetooth accessories and connect and disconnect to ones you find. That we know. It also says you can vend services from your app, turning the iOS device into a peripheral for other Bluetooth devices. And it says you can broadcast your own iBeacon information from the iOS device.

Leo: That was, to me, what was fascinating is that Apple stealthily had included this in all iOS devices since the iPhone 4s. I mean, this is hundreds of millions of iOS devices. Now, it has to be enabled. It's an API. I mean, it's not like all of a sudden your iPhone's going to say hello, hello, hello.

Steve: And so what Forbes is proposing is that what will follow iBeacon is iWallet, and that Apple will end up producing an immediately pervasive electronic payment system based on...

Leo: Sure. Walk into a store. You'd have to opt-in. You have to turn it on. You'll have to run an app. But you could walk into a store, and they go, "Here's your coffee. Thanks. See you, Leo." And it's done. The transaction is done. I think that's - now, remember these iPhone 5s's have fingerprint readers. They could tie it to that for authentication.

Steve: Yup, yup.

Leo: Are you worried? I mean, what is the implication of this? Isn't it possible to lock it down if it's implemented properly?

Steve: I guess I - yes. I would say, as long as the user receives something on their

screen, and they look at it, and it's not spoofable, and it can't be intercepted, and they affirmatively acknowledge it, and...

Leo: I mean, you could just do a PIN. But I think the point of all this is friction-free.

Steve: Yeah, you want low friction.

Leo: You don't want to have to make the user do anything. They walk in, they take their coffee, and they walk out.

Steve: Yeah. Right now with our iPhones, of course, at Starbucks you show the scanner the barcode on your Starbucks app on your iPhone, and that debits it from your account. So I like the fact that there's that kind of, that level of interaction. Radio is a little more spooky. I mean, and we've seen it being exploited. So we'll see. I mean, if nothing else, maybe a privacy concern because somebody else can be monitoring this. We'll have to - I've not looked closely at the spec and what they're doing. But I agree with you it's going to be popular.

Leo: Oh, yeah.

Steve: Real quick, I ran across a list of two-factor authentication websites. I just wanted to share it with our listeners because it's kind of cool to see it. I created one on my short bit.ly links: bit.ly/2falist. 2fa is two-factor authentication. So 2falist, all lowercase. That'll bounce you over to Evan Hahn's site, where he's just, as his own little project, been maintaining a list of major popular sites that support two-factor authentication right now. And we can assume that will be growing over time. So sort of disturbingly short list. I don't know if it's extensive. Maybe if anyone knows of any others they could drop him a piece of email and say, hey, add this. But certainly the ones that he's got are ones we know about and that we've talked about on the podcast. So it's nice to have that, just sort of to browse through.

Leo: Yes. But you're right, it is. It's kind of surprisingly short.

Steve: It's not that long.

Leo: But you know what, it's good. He has links to enable it underneath each. So you can go through and see which ones of these services you use, and at least if they have it you can enable it from there, which I think is great.

Steve: Yeah. Okay. So now miscellany things. I'll run through these quickly because it would be fun to get to some questions. I have been experiencing failure of Touch ID. And it's been annoying me for the last couple weeks. And I was so glad then to run across many other reports. It's happening to many people. And when I tweeted about it, I got lots of followers who said, oh, yeah, me, too, me, too, me, too. So it's a phenomenon. One of the things that I noticed is that it seems to be, for me, temperature sensitive. On

a cold day, where the phone is cold and my hands are cold, it's much lower recognition level than when everybody's warmed up.

So I had a couple ideas, and then I've seen them elsewhere. One was, if in fact you don't get recognition when your finger is cold, record/train Touch ID in a different slot. Because you've got five, train the same finger in a different slot under the cold conditions in order to have it recognize that. Or, if nothing else, just do multiple trainings of the same finger in multiple slots to give it more samples. And other people have found that maybe the polar rotation orientation isn't working as well, and so to deliberately record a straight-up finger and record a 90-degree angle finger one way and then a negative 90-degree angle finger, again, all in separate slots.

Leo: If you retrain, like you just delete that fingerprint and retrain, it gets better; right? That fixes it.

Steve: Yes, it works again, yes.

Leo: So strange.

Steve: So I'm calling it "finger fade."

Leo: Finger fade.

Steve: Apple iTouch Finger Fade because - and, see, I would - I don't know what the logic is. Were I designing this, and I decided that the finger I had just seen was the finger I knew, but it was also giving me new information, I might be trying to mature my knowledge of that finger over time. That is, I recognize enough that I believe it. And, oh, look, I'm getting some more on the side here that I hadn't seen. Let's add that. So if there is an algorithm, where they're hoping to further evolve their finger ID, maybe it's not working right. Maybe it's like that that evolution is going sour. Who knows.

Leo: It's not possible your fingerprint is changing over time. Dr. Mom says this is - because they use this in hospitals a lot for access to meds and things like that. She says this is not an uncommon thing with all fingerprint readers. I don't have enough experience to know that.

Steve: We know that fingerprints themselves don't change over time because you have a huge...

Leo: Not substantively.

Steve: Right. We know that you can take fingerprints when you're young, and they're still valid when you're old. That's Forensics 101 of law enforcement is your fingerprints are your fingerprints are your fingerprints over time. What we're hoping, I mean, all of the hype said that the capacitive technology was ignoring surface dirt and was like

reaching into and looking at the meat of your finger. I just think training it under different conditions is probably a good thing. And filling up all five slots with the same finger, you know, why settle for one?

Oh, there was someone, I did read that, if you fill up all five slots, then recognition takes longer. And that certainly makes sense that it would have to, like, if it didn't find it on the first one, then try the second one and so forth. And so if it's going to be unhappy, it's going to take longer to be unhappy if it's got more cases. So anyway. Also, real quick is what looks like an actually usable keyboard for the iPhone.

Leo: Oh, no. Not Ashton Kutcher's keyboard. You're going to get this?

Steve: It's on - no, it's Ryan Seacrest's.

Leo: I'm sorry, Ryan. I confuse the two. They're equally vapid. Really. You think this is good, huh? It looks just like...

Steve: Look at it. It looks...

Leo: ...a BlackBerry keyboard.

Steve: That's, hello, yes. I mean, my BlackBerry is on the sidelines right now being sad. If I could have the BlackBerry keyboard on the bottom of my phone - the other thing this does, and I thought that it was an interesting point, is as is the case when you have any keyboard attached to your iOS devices, the software keyboard no longer deploys. So you get...

Leo: Which is good. Much more screen room.

Steve: Yes, much more screen real estate because in some cases, in some apps and uses of the iPhone, it just scrunches your screen down to nothing because the keyboard takes up so much space. Anyway, this is TypoKeyboards.com. Anyone who's interested, TypoKeyboards.com. It's a case. So something comes down from the top and plugs in at the bottom. You still get your Lightning connector on the side. It uses Bluetooth to talk to the phone. Maybe, I'm not sure. Maybe, if it's going to plug into the Lightning, in fact, I'm not even sure that uses Bluetooth. The screen shot shows Bluetooth on. But iOS keeps turning it on every time you update Bluetooth, which is annoying. So it sort of comes in from both sides. And anyway, it's supposed to be middle of next month, middle of January. So I'm jazzed. I will have one. And I will report.

Leo: How retro.

Steve: Because, boy, if I could have a real keyboard on my iPhone - because I'll tell you what I've found, Leo?

Leo: What happens to the Home key? Because it covers the Home key.

Steve: Yeah, it does. Good question.

Leo: I guess the...

Steve: Must be there somewhere. Oh, there it is. I see it in the far lower right corner.

Leo: Next to the - yeah.

Steve: That looks like a Home key. If you look at the big picture in the show notes. So, oh, maybe. That would be cool. What I have found is, of all the keyboards - iPhone Portrait, iPhone Landscape, iPad Portrait, iPad Landscape - my very favorite is the mini in portrait. So the mini, the iPad mini, held upright in one hand, that keyboard is just the right size for me to type on.

Leo: Well, you can thumb type, too, because it's narrow enough you can hold it like you would a BlackBerry and thumb type. And I presume that's why you like it.

Steve: Yeah. And splitting the keyboard doesn't work for me, either. I'd like it all together. So anyway, maybe we're going to get a cool keyboard for the phone. I wanted to let people know. Also I have found and am loving something called FocusAtWill.com. It is curated music for people who want to work. And so it is much like the Liquid Mind stuff. It is $35 a year, but you can try it for free. For free you get a 300-minute or five-hour loop of the same music, so after a while it's going to be repetitive. But someone tweeted me about it. I tweeted it out after using it for a day and loving it. And I've had a ton of feedback from my followers who have tried it and are completely addicted.

You can choose between Classical, what they call Focus Spa, Up Tempo, Alpha Chill, Acoustical, Cinematic, Ambient, and I like ADHD Beta Test, whatever that is. I haven't been moved to go there yet. And then in each of those you choose low, medium, and high intensity. So you have a wide matrix of music. And it's just really good. I mean, it's completely pulled me off Pandora and everything else. I really like it. And for free, you can try it for free. There are iOS and Android apps. It'll also run in any browser. So you can just - a browser will play this, and you can see what you think.

So it was Leif Jantzen who tweeted this. And so thank you, Leif. It's changed what I do. And lots of people are saying they love it. A little bit cheaper than Pandora. And, no, not nearly as much flexibility. But if this is what you want, it really does deliver it.

Leo: I'm going to play a little bit of it right now. Just relax. This is the Classical. Actually, nothing's coming out. I don't know if I've done something wrong. Seems to be playing, but I don't hear it. Oh, I have the wrong...

[Music]

Leo: So this is actual classical music.

Steve: Yeah. But it's not like huge tympanis going bang, bang, bang.

Leo: There's no boom, boom, bada boom.

Steve: There's never any lyrics, but...

Leo: What do you listen to? You like Alpha Chill, Acoustical, Cinematic?

Steve: Alpha Chill is nice, so try that.

Leo: All right. Let's see some Alpha Chill. I like - you know what, I might get this because I do like instrumental music in the background. I prefer classical, myself, but...

Steve: And I've been listening to Classical. It's really nice. Also Acoustical. Try Acoustical.

Leo: Yeah, I like acoustic music, actually. Whoops.

Steve: It'll be a little piano and something, a little...

Leo: By the way, what we're listening to is free. Right?

Steve: Yes.

Leo: Yeah. So I don't know what the limits are on the free version, but...

Steve: Oh, I do. There are no limits except it's just a five-hour loop. So after five hours...

Leo: Oh, it repeats. Yeah, yeah, you said that, yeah, that's not bad. Five hours is enough for anyone.

Steve: So listen to this.

Leo: This is beautiful. And you use this for programming and writing and stuff you need to...

Steve: Yes. When I'm writing or when I'm coding, and when I want to block out background noise and just zone in. And, I mean, there's all this bull [bleep], oh, excuse me. There's all this...

Leo: I can see this music has really relaxed you, Steve.

Steve: There's all this other stuff they've got where it's like, all the science of psychoacoustics and...

Leo: Okay. I want to see the ADD stuff. Let me see. ADHD Beta Test. So this is the AD - oh, this is like trying to - trying to...

Steve: Oh, my lord.

Leo: I'm ADHD. I need other stimuli to keep my brain - see, the whole thing with ADHD, you wouldn't know this, is that we actually, those of us who have this, our frontal lobe is under-functioning.

Steve: I guess I don't have that.

Leo: You don't have it.

Steve: That would just make me nauseous.

Leo: And so the reason people are ADHD is they need to stimulate their front lobe. And if they can keep their frontal lobe stimulated, then their mind works like yours does. So this, I guess the idea is stimulate the frontal lobe. I'm a little over-stimulated, frankly.

Steve: Just shoot me now.

Leo: [Laughing] Wow, that's Focus...

Steve: FocusAtWill.

Leo: ...AtWill.com.

Steve: And also, again, iOS and Android. So you can put it on your phone. You can put it on your Pad. And it's just there. It's just, I mean, I've just been having a fabulous experience with it. So I know that not everyone follows me, or some who do follow so many people my note of it might have been lost. So I wanted to let our people know, our listeners, that it really looks good.

Leo: Hmm.

Steve: I found, after we talked about our faster-than-light drive, the warp drive last week, there is a one-hour video presentation that Dr. Harold White did following his paper, that we had the PDF last week. There is a YouTube presentation. And a couple people, first of all, we've provided entertainment for a subset of our listeners who apparently were laughing so hard they fell off their chair. But that's fine, too. And one person tweeted, saying, he said: "More like less impossible. The solution described essentially requires several tons of exotic matter having negative mass." So okay, Leo, maybe you were right, that it won't be anytime soon.

Leo: Oh, well.

Steve: And someone else did correct me. We talked about going to Alpha Centauri in two weeks, which actually just came from text that I read about it. It turns out it's 0.43 years to Alpha Centauri. So that's 157 days, or 22 weeks. So maybe someone dropped a "2," and they went from 2 to 22. It's 22 weeks. But even that's practical, if you can find a couple tons of exotic matter. So maybe we can get the Large Hadron Collider to create the exotic matter for us. Then we'll stick it into our warp drive spaceship, and off we go. And really, a couple good Hamilton books will get you to Alpha Centauri. That's really all you need.

Leo: That's how you measure time passing now?

Steve: Exactly.

Leo: That's 4.3 Hamiltons.

Steve: That's 4.3 Hamiltons [laughing].

Leo: I like it.

Steve: Also, a quick sci-fi update. "Almost Human," that I talked about on Fox, I'm getting a ton of positive feedback from people who were thank you, thank you, thank you for putting me onto it. I got a note from someone who tweets as @weckman, who said: "SGgrc, did you know Fox is 'pulling a Firefly'?" Now, the moment I saw that, I thought, oh, no, no, no, don't cancel this. And he said - but he says: "They've released the episodes out of order for unknown reasons." Apparently what they've released is 1, 5, 6, 7, 8, 3. And I thought, huh. And Monday's was really good. And so it motivated me to

again remind everybody that "Almost Human" on Fox is, I mean, it is - I would say it's the quality of "Firefly." And we haven't seen anything like that for a long time. And of course everyone's nervous because Fox killed "Firefly" for, like, I don't remember what Joss said was the reason. But let's hope that "Almost Human" survives. I'm really liking it. It's the best thing I've seen for a long time.

Leo: I watched the first episode. I did enjoy it. I was a little disappointed because the first android he got was such a dud. And I thought, is this really the show? And then I realized, oh, no, that's not the android he's going to get.

Steve: Oh, and last night's, if you want to just jumpstart, you could watch last night's out of sequence. Apparently it was the third one made. It's the best so far. More character development. Just I'm impressed by the writing. They did extra things they didn't have to do that you don't normally see. Anyway, it's definitely on my must-see list. And people have been reminding me about "Continuum." I saw the first episode, and it's like, oh, maybe it was just a bad day for me. I'm going to go back, now that there's a bunch of them, and maybe run through them and see what I think. I do have a SQRL update...

Leo: Squirrel!

Steve: ...that is significant.

Leo: Okay.

Steve: Yes, thank you. A major milestone yesterday. I posted the final piece of the protocol, the so-called semantics. I had put up the syntax before, which is the way the endpoints will communicate. The semantics, of course, is what they will communicate. That's now online. The denizens of the SQRL newsgroup at GRC are plowing into it, figuring out what I've done, and I'll be looking at their feedback. This was just yesterday, so it just happened.

What this means, though, is once the dust settles from this, I finally start writing code. There is enough there that everyone, once we agree upon this, will be able to write code to implement SQRL. So I'm very excited. That generally goes pretty rapidly, compared to just, I mean, it's been a real process getting to where we are. But we're there now. We have a robust specification, future oriented, open ended so it can grow and do other things. And I'm ready to write code. So the people waiting for v6.1 of SpinRite will be glad of that also because it certainly represents a good step in that direction.

And speaking of SpinRite, actually in this case what we've got now is 6.0. In keeping with the spirit of a Q&A, I ran across a question in the mailbag from Jared in Australia, who asked a question about SpinRite and solid-state media. He was a little bit confused about SSDs because he said: "Hearing a previous question on recovering data on SSD, I don't really get this. On the one hand you say yes, SpinRite can be run on SSD, and you don't do anything to prevent this. But you also say that SSD and flash media have limited write cycles. So running SpinRite on SSD also wears them out. How can this be a good thing?"

And so I wanted to, for people who haven't heard, and also for Jared, to answer his

question, SpinRite has what we call Level 2, which is a faster operating, read-only scan which looks for any problems on media. That makes much more sense, I mean, well, it's the only thing that makes sense for SSD, but it also makes sense because SSDs don't have defects the way hard drives do, that is, physical defects, where it makes sense to pattern test in order to find them. So Level 4 is a read-and-write test that you do not want to run on SSD because it would tend to fatigue the SSD substrate, but which you do want to run on drives when it makes sense because it's like sort of a deeper level of testing.

So, yes, SpinRite runs on SSDs all the time. And in fact, in the show notes I posted something else, not from Jared, but Leo, you can show this. It's a little frightening. It's actually from a customer. It's a customer photo that was supplied showing SpinRite running on a brand new Kingston SSD, where the SMART system is already getting concerned. There's a red dot at the far right end of that bar because what's happened is essentially SpinRite is running, and the SSD is using error correction to such a high degree that it...

Leo: Look at the error count. It's 2.9 million.

Steve: I know. It is horrifying. And so this should forever disabuse us of the notion that, because it is solid state, it is error free. What we're seeing here is perfect evidence that they have pushed the density of SSD storage, that is, commercial pressure to cram more density essentially in smaller and smaller cells has resulted in solid-state storage having an effective error rate, arguably right up there with hard drives. And not only is the absolute error count horrifying, nearly 3 million, depending upon how far he went.

But the other thing that is a concern is look at the spread between minimum and maximum. What I do is those numbers are errors, corrected sectors per megabyte. That is, I deliberately do the math so that the units have meaning. The unit is corrections per megabyte. And so what this says is that there was a region on the SSD where we were only needing to do about 1,500, the minimum, 1,500 corrections per megabyte. But there was a different region where we were having to do about 25,000 corrections per megabyte, meaning there's, like, a bad area on this SSD compared to, like, the better area. What you'd like to see with something semiconductor would be some uniformity across the surface.

So I'm seeing instances now where we're seeing the quality of these drives beginning to come down. And I've never mentioned it on the podcast, and I sort of meant to a while ago, but our listeners may have noticed that warranty length kind of quietly crept downwards. It used to be that I think they had, like, three-year warranties. They've dropped them to two years and one year, just sort of very quietly saying, eh, we don't want these back after a year, so you're on your own. Because warranty periods used to be substantially longer, and that's not the case any longer.

Leo: Do you recommend against SSDs? Or should you go with business class or lower density?

Steve: I'll tell you what I would do. I would do what Compaq used to do, Leo, at this point. Compaq used SpinRite on their loading dock to prequalify their drives. The manufacturers didn't like it, but they had no choice. Compaq over-ordered what they wanted. They ran SpinRite, and they returned the worst of those.

Leo: Isn't that interesting.

Steve: I would buy two hard drives and run SpinRite on them both and see if they're different. And if they're different, return the weakest one.

Leo: You are the weakest link.

Steve: Why not? Why not? You know? Everybody will take them back without questions.

Leo: It's a great idea.

Steve: Yeah.

Leo: Buy two, return one.

Steve: I would buy two.

Leo: Buy three.

Steve: Well, see, and that's the other thing, is that people sometimes show me a screen like this and say, is this bad? And it's like, I don't know because every...

Leo: What's normal?

Steve: Every make and model differs. Now, what is definitely bad is this one. That SMART error, that is the drive itself rating itself as weak. We're, I mean, if the drive is saying, wow, I'm doing more error correction than my firmware controller thinks I should be. Because the firmware controller in the drive is separate from the medium. The medium is like chips out there. The firmware controller has been designed to, like, it's counting error corrections, and it's saying ow. I'm showing - I'm dropping the health rating because this is more ECC than I expect. So that's an absolute reason to send this back.

And this was a brand new Kingston SSD that the guy bought. And he said, this doesn't make me feel comfortable. And I said, and it should not. But what you can also do, the reason - when people just show me a screen, it's like I don't know if that error count is bad or good. But if you had two drives of the same make and model, if you bought a pair, by comparing them, if one is generating a lot more errors than the other, that's the one you don't want to keep.

Leo: And I guess if you bought two or three, if you bought three you'd have...

Steve: Then you'd really know.

Leo: You'd really know.

Steve: There'd be a, yes, you'd have two that are similar and an oddball that's like, ooh, boy, you go home.

Leo: Yeah. Cool. Cool. All right. Are you ready for questions, Steve?

Steve: Well, now, okay. We are, but we're approaching a hundred minutes on the podcast.

Leo: Is that your new benchmark, the 100 minutes?

Steve: Well, I'd just - I always put the minutes down.

Leo: It's a good number.

Steve: Yeah, it seems, I mean, that's, like, that's a good podcast.

Leo: Yeah.

Steve: So what we'll do is we'll do a couple questions, and then we'll just continue next week with another attempted Q&A.

Leo: We've been attempting this for a month now. Here we go, listener-driven potpourri. Doug, whoa, let me just...

Steve: Gernz, Gernetz...

Leo: Germetzky, I think.

Steve: Gernetzky, yeah, Doug Gernetzky.

Leo: Nice Polish name. Doug Gernetzky in Appleton, Wisconsin, he caught Steve's attention with his subject line: The 800-pound gorilla you won't talk about. HealthCare.gov just had an investigation on web security. I heard about it all week and was surprised not to hear about it on Security Now!. If you didn't cover it because of Leo's politics [laughing]...

Steve: I know. Keep going.

Leo: Oh, lord - will you discuss it elsewhere? It's kind of a big deal, after all. It's only the medical records and personal info and passwords for America. Thanks, Steve. Phillip Lane also asks: Have I missed it? I was going through show notes for past shows looking for an informative show on HealthCare.gov, and I see nothing in the notes. You've not even touched on the subject? There's so much information out there on the ugliness of this rollout, capacity and security issues, I find it hard to believe you haven't covered this news. Well, it's because of me. I'm a lefty, and Steve doesn't want to go against me.

Steve: It's actually not because of you at all.

Leo: Oh.

Steve: Phillip Lane says: Just looking for some nonpartisan reality. And for one thing, the whole issue is so fundamentally partisan...

Leo: Highly partisan, thank you.

Steve: ...that it's difficult to know what reality is. But mostly I just assume that it is the nightmare catastrophe, catastrophic disaster from hell. I mean, it's like, why bother talking about it? Just assume that it is a massive fiasco. Apparently half, about half of the state-run sites do what we were talking about earlier. They log you in briefly with SSL, and then you have a persistent cookie which allows anyone to take over your session and have access to all of your data. I mean, it's like - I was explaining this to Jen because initially it wasn't even working. I mean, like two people could use it at once. And I just, I mean, I didn't assume anything different. It was like, okay, this is a fabulous fiasco disaster on the part of the government. And when I heard that it was loading, the website was loading 60, six zero, different individual JavaScript files from all over the place in order to function, it just, okay, you know, forget about it. So...

Leo: The problem is it's hard to get real information. You're right, it's certainly assumable that there's a problem. But CBS News, for instance, aired a report about security issues that was wrong.

Steve: Yeah, see, none of them know what they're talking about.

Leo: That's the real problem.

Steve: That's one problem. And I don't know what I'm talking about. I only know how hard it is. I mean, everyone who listens to this podcast knows it is so difficult for old-school serious Internet companies like Amazon or Twitter or Facebook, I mean, full of serious techie gurus, to get this right. It is really hard. And this was the point I made to Jen was, I said, Jenny, it doesn't even work. I said, so, I mean, working is the first bar

you have to get over before - and unfortunately security is always an afterthought. I mean, we see this over and over. I mean, we should just assume that there is no security whatsoever, and that some year after it finally works, maybe it'll start being secure. I mean, it's such a catastrophe, I've just ignored it. So the reason I haven't talked about it is that I really had nothing definitive to say. I need details. I need specifics. And it's just like...

Leo: Fox News, CNBC, CBS, you're not going to get the technical details you need.

Steve: Is the ocean wet? Well, yes. You know? It's like beyond security. It's just, I mean, it would just hurt you to start looking at it. So just, I mean, there's just nothing that I have to say. I don't want to look at it. I don't want to take my own time. People would much rather have SpinRite 6.1 and SQRL than me going and looking at HealthCare.gov and taking it apart. I mean, just assume...

Leo: Well, it's funny because of the four security experts that testified in front of Congress, one said exactly that Avi Rubin, who we know as one of the great security guys of all time, is a professor at Johns Hopkins. He said: "I would need to know whether there are inherent flaws versus superficial problems that can be fixed. If they can be fixed, that's better than shutting it down." He, like you, said I need to - we aren't getting the information to make any conclusion. Yes, there's possibly cross-site scripting errors. It's certainly possible to phish people and trick them. And some of this mainstream coverage is simply that. You could get a link in an email that says it's HealthCare.gov, and it looks like HealthCare.gov, but it's not. But that's going to happen with any site. So it's hard to really know. And I think that...

Steve: Well, see, and that's my point. It's uninteresting to me. It is such a disaster that it's uninteresting. It's just like, okay, it's horrible. And I'm not blaming anyone. To me its horror is nonpartisan. Its horror is a fact. And someday they'll get it fixed. I mean, now they understand that there's problems. Security is like the next thing they're looking at after getting it working. It still has a very pathetic number of total - I just saw the number this morning, twice in November as in October, but we're still down at a quarter million, which is way shy of seven million that they were, like, was some benchmark they were hoping for. So, yeah, I mean, I just - it's uninteresting. I mean, what's interesting is something I can get my teeth into. This is just, you know, it is the 800-pound gorilla, and I don't want to try to bite it.

Leo: Well, and there's a lot of non-, I mean, a lot of partisan information. The four security experts who testified for Congress, one of them is the former senior law enforcement advisor to the Republican National Convention and a cybersecurity analyst for Fox News. I have a feeling he has a partisan point of view. A lot of these guys are, in fact, not what I would call - one guy's a former CSO of Diebold, the famous electronic voting company that had the worst security record of anybody I've ever heard of.

Steve: In this situation, I would say believe the worst that you hear.

Leo: Assume the worst, yeah.

Steve: Yes, because we have - with security we have to assume the worst. It is, as we know, it's about the weakest link in the chain. I can't even begin to imagine the chain that this thing has. I mean...

Leo: The problem is a lot of people have to use this. Their employers are saying, hey, we're not going to give you healthcare or whatever. I mean, the security advice would be just don't use it till they fix it. But a lot of people have to.

Steve: Well, I mean, if there's ever been a need for making up a new password when you identify yourself...

Leo: Well, that one, for sure.

Steve: Do it here.

Leo: But I presume you've got to give them your social. I mean, you're giving them a lot of information. I don't know if you have to give them your social, but you've got to give them a lot of personal information, including...

Steve: Yeah, no, I mean, it's a catastrophe. It's a disaster. So I hope I've satisfied any listeners who feel like I was afraid to talk about it. It's not at all that I've been afraid to talk about it. It's like, where do you start? I mean, we'd like to talk about something else. But if we have to talk about that, it's like this will become the HealthCare.gov disaster from hell podcast.

Leo: Yeah, assume the worst.

Steve: I mean, oh, my god, yes.

Leo: I like it. Yeah, it's not good. It's not good.

Steve: It's beyond - it's like it doesn't even work. So of course it's not secure. It has to work before it has a hope or a prayer of being secure.

Leo: Yeah, it doesn't even work, yeah.

Steve: And someday they'll do that.

Leo: Yeah. And by the way, Web6121's pointing out that you don't have to use the website to sign up. You can phone in, which is probably a good idea.

Steve: The problem is...

Leo: We don't know how well secured the databases are.

Steve: I was going to say the problem is they're probably using the website on your behalf.

Leo: Maybe they are, I don't know.

Steve: When you phone up, yeah.

Leo: We don't know. That's kind of the bottom line. Ron Bogner in Glendale, Arizona has some thoughts about 256-bit identifiers. Ron says: It amazes me that people who have a problem with 256-bit identifiers do not mind a 10-character username paired with an eight-character password. And he gives a sample BitTorrent Sync secret, which is a very long, 256-bit identifier. Here's a sample username and password: johndoe@yahoo.com, Monkey12. Obviously the 256-bit secret is magnitudes more secure than the password/username combination. And that would be true even if you used a random 15-character password instead of something more typical, like Monkey12. He does some math. We've said this over and over again. I think the issue isn't so much that we don't think it's secure, the fear of collisions.

Steve: Yes. Yes. And in fact that was my math that I added there.

Leo: Oh, that's your math, oh, all right.

Steve: Yeah. It was just because I wanted to point out, people assume that 256 bits uses 8 bytes. But if you use the full character set, a byte is 256 characters. Well, first of all, we immediately throw a bit away because we have 127 characters. But then when you throw all of the - and I'm sorry, 128, technically, although null is invalid, so 127. And so that would be 7 bits. So you would divide 256 by 7 in order to get the number of characters, except that many of those are control characters. And I went over to my own Password Haystacks page, just to remember what the alphabet size was. And if you use all upper, all lowercase alpha, all 10 digits, there are 33 typeable special characters. So that's a total of a 95-character character set.

And one of the neat tricks I've always loved, if you wanted to know, like, how many bits 95 was, because it's an odd number; right? It's bigger than 64, which is 6 bits, but it's smaller than 128, which is 7 bits. It's 95. It's kind of somewhere in the middle. Well, it's like, okay, where? So if you take the log of 95 divided by the log of 2, because 2 as in binary, that ratio gives you the exact number of bits, binary equivalent because you've

dividing by the log of 2, that 95 is. So it turns out it's 6.5699 bits per character. We'll round it up to 6.57 bits per character. Then you can divide that by 256 bits, which is the actual amount of entropy possible, and that tells you that it's about 39, it's 38.966, or 39 characters, rather than, for example, just 32. So you get actually a bunch more characters when you treat the alphabet correctly.

But he's right. Certainly that's unique. And I think, as you said, Leo, what people are concerned about is collision. And we do have a question we'll probably see next week, further on down here, that talks about that, how one listener thought we were being a little too glib when we said it's impossible. Because he says, well, it's not impossible. It's like, okay, Spock. So...

Leo: I believe your logic is flawed.

Steve: So, but that does mean that we have to be very careful with randomness because we know that that's a problem. And I think the other thing that makes people uncomfortable is that, when you are creating an account somewhere, you put in a username, and it says, oh, that username is taken. So you go, oh, okay.

Well, unfortunately, in the process of being told that, you now know somebody else's valid username, and you could then start guessing their password. And of course the reason that email is so handy is it is known to be unique. But the only reason it's known to be unique and for you is that you already - somebody else told you that there was already a johndoe@yahoo.com. That's why you had to go johndoe2048 or something to get a unique version of johndoe. So the idea is that then you reuse that everywhere because that's been assigned to you. And that's the way we avoid collisions.

So the fact is, except for the Spock instance, with really, really high-quality random numbers, or pseudorandom numbers, there is no practical chance of a collision. There is much higher chance of anything else in the universe going wrong than there being a collision of two identities. I mean, your computer could fail. Your power could fail. We talked about the asteroid hitting the Earth. I mean, everything else is going to happen before there's going to be a collision, if we choose random numbers carefully, and they're 256-bits long.

Leo: So we are out of time. But I want to read this 11th question because I think that it would be good to address this now. And then we'll put some more questions together for next week.

Steve: Well, I'm going to carry these on to next week.

Leo: Carry these on, yeah.

Steve: We'll try to get to them.

Leo: Scott Schramm in Philadelphia says what about show notes? I love the Security Now! podcast. Great topics. Usually after the show I go to the TWiT page

[TWiT.tv/sn] in search of the links talked about in the show, but the show notes are always nonexistent. There is a link to show notes, but it's always blank. Seems to be the same with other TWiT shows. I cannot find any show notes on the GRC page, either. The text transcripts are great, but I don't have time to read that. I'm looking for a short, bulleted list of all the topics mentioned and any links included with those topics. Can you please ask Leo about this. Thanks.

Steve: Okay. So we need to explain a little bit about how we operate here. You and I get together for two hours and do this, and then sort of go our separate ways. After a few hours your guys put the audio together, I download it and then shrink it and get it off to Elaine. And then she produces the transcripts, which come back to me. For the last three weeks I have, counting this one as No. 3, I have been, finally, for the first time ever, posting the show notes on GRC.

So if you go to GRC.com/sn, which is the main Security Now! page, you will see now an additional icon in the lineup. There used to just be five: high-quality audio, lower quality audio, and then the transcript in three forms (text, html, and PDF). There's now a sixth one, and it's the third one over. It's after the two audios. Looks like it was one of Microsoft's old icons, it's a little cubes tumbling together. That is now a PDF of the show notes that you and I go through at the beginning of the show. I tweet it before the show because I've had a lot of positive feedback from our real-time, our live listeners, saying, hey, it's so fun to be able to follow along, which everyone can now do.

So if your guys, Leo, want to do something with the show notes, they're welcome to. But for what it's worth, I will from now on always have them on GRC.com.

Leo: Yeah. And so, and this is a terrible thing about everything that we do, is that we are highly resource constrained. It probably feels to people like, and we certainly encourage that, that this is a big operation, and we've got editors and big studios and stuff. But one thing we don't really do a very good job of, we had hoped to crowd-source show notes, and we don't do a very good job of that. For a while we did do a pretty good job of keeping the Security Now! show notes up to date. I would paste in your notes because you do very good notes. Of all the hosts, I think you, Paul Thurrott, and Mary Jo Foley do the most complete notes. And it would be - and I should be just putting them on the wiki.

So that's why it's a wiki, by the way. It was our hope that people would, who listened to the show regularly, would go on the wiki and parse those notes and so forth because I don't have time to do that. And I really - I haven't even had much time to - I'm busy doing the show, so I don't have time to paste it into the wiki, and I apologize. I slacked off on that. We should probably have a full-time person doing that. I just can't afford it. And it would be probably two full-time positions. So we're talking 100,000 a year.

Steve: And in fact, yeah, and if you were going to do notes for the other podcasts that didn't, like, generate their own notes in order to drive them...

Leo: Well, that's what I mean, it would be - yeah, right, be a lot of work.

Steve: ...that would be massively, massively labor intensive, yeah.

Leo: We have always rundowns somewhere, you know, we share rundowns and Google Docs and stuff. And we do our best to put together notes. But I understand it's a complete failure. We've never done a very good job. And I apologize. Thank goodness Steve is putting his show notes up. And that's everything you want. So to answer that question, it's getting done there. But...

Steve: They are always available from now on there.

Leo: It is a global failure of mine. And I don't know what to do about it.

Steve: Well, it's just maybe there will be something, but at least you can always find them for Security Now! at my site.

Leo: Yeah, thank you. And I will talk to Lisa and see if we can figure this out. You have to understand, hiring a staff person to do this is very expensive. It's not just the salary, it's the benefits, it's all of that.

Steve: Well, and I can't speak for the value of the notes. I guess I'm liking that I'm now publishing them because what I used to have to do before is I was, like, tweeting these links all the time. And I would be saying, oh, I just tweeted this link, I just tweeted that link, to help people who want to, to go find it. Now, this is what drives the podcast. So the notes are there. Anyone who wants to find a link knows where they're going to be. So, but my question would be, do notes really make sense for all of your other podcasts? Are they going to have stuff in them that people are going to want to find and follow up on? I don't know.

Leo: Yeah. Well, show notes are really important, primarily because it makes it easier for a listener, but also it makes it searchable. And that's where your transcriptions are so great, and it's one of the reasons we're going to start doing transcriptions on more shows. We'll talk about that on Inside TWiT later today. We're going to pay the money to do that, anyway. That's not show notes, though, as Scott pointed out. That's too voluminous. So we've got to figure it out. One of these days. There's so many things I'd like to do, Steve. You know how much it's going to cost. We're redesigning the TWiT.tv's website, which we redesigned last year, and it's not working. That was $150,000. It's now almost twice that to do it again. It's expensive. I wish I had more money. Ahem. But at least we've got our fine sponsors. At least we've got this great show. And I thank you all for being here. And I especially thank you, Steve.

We do this every Wednesday, at least for the next week, two weeks.

Steve: Yeah, one more. Wait.

Leo: Two more.

Steve: Oh, yeah, because Christmas will also be on Wednesday.

Leo: Yeah, 11:00 a.m. Pacific, 2:00 p.m. Eastern time, 19:00 UTC. Then we move to Tuesdays at 1:00 p.m. Pacific, 4:00 p.m. Eastern time, 21:00 UTC. That'll start January 7th. And Christmas...

Steve: I'm going to like that.

Leo: Go ahead.

Steve: I was going to say I'm going to like that because it'll give me another two hours to prep. I worked from 5:00 a.m. until 6:00 to get this ready.

Leo: Yeah. Now you can get up at 7:00.

Steve: Yeah. Well, or just be a little less frantic in the morning.

Leo: Anyway, I really appreciate your doing that. And I understand that it's perhaps for some a deal breaker, but we do make on-demand audio and video always available, and you can listen at any time, at your convenience, after the fact. And it's got everything we talk about except for the bad words, which are bleeped, at TWiT.tv/sn. Steve has 16-bit versions at his GRC website, GRC.com. He also has a little thing called SpinRite you ought to have, if you've got a hard drive, even a solid-state one - I like this idea. Buy three.

Steve: I know.

Leo: SpinRite all three, keep one.

Steve: Yup. Exactly.

Leo: I like that idea.

Steve: Find the best of the crew and send the other two back.

Leo: You could do that, too, at GRC.com. And if you have more questions for feedback, go to GRC.com/feedback. That's easy to remember. Thanks, Steve. We'll

see you next week on Security Now!.

Steve: For Q&A continued. Thanks, everybody.

Leo: A little more.