Listener Feedback #126

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-318.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-318-lq.mp3

Leo Laporte: This is Security Now! with Steve Gibson, Episode 318, recorded September 14, 2011: Your questions, Steve's answers, #126.

It's time for Security Now!, the show that covers protecting you and your loved ones online, privacy, security and all that. And who better to do Security Now! than my buddy right here, Mr. Steve Gibson of GRC.com...

Steve Gibson: Yay.

Leo: ...the Gibson Research Corporation. Yay. And the crowd goes wild. Steve is the creator of so many useful programs to protect you online. And now we've got 318 Security Now! shows that you can study up on, which is great. Welcome, Steve.

Steve: Hey, Leo.

Leo: Q&A today.

Steve: So we have a Q&A. Not much, happily not much news. Boy, we've been buried in security news the last few weeks. So very light on the news this week, but we've got some on the long side questions from listeners, but that I felt were important. So you'll be doing some reading today, Leo. And but I think in all it will balance. And I have a very exciting announcement of a new science fiction author discovery to share with our listeners. So lots of good stuff, and a super-short SpinRite mention, just to keep that on everybody's radar.

Leo: Have you been following, have you been watching the Microsoft Build Conference, the new Windows 8? I guess you of all people probably care the least about...

Steve: Leo, from everything I've seen, just shoot me now. I mean, I look over, and my Windows 7 that I built just to run Skype on is now telling me that XP has 936 days remaining. And I just think, good, thank god, or thank Microsoft. Because the next thing I'm going to build, I'm going to update my processors to i7s, but I'm going to use 64-bit XP. I cannot stand Windows 7. I've tried to get used to it. It's on a few machines. It just - every time they do another iteration, they move it further away from something that just makes sense. And Elaine had to switch to Windows 7, oh, my goodness, you wouldn't believe how, I mean, I could relate to everything she was saying about how she just hated it. And so I'm so glad I've got three years left of XP support. And then I look at this Windows 8, and it's like they've just gone off the deep end. And I meant to ask you, what does Paul think about this?

Leo: Well, I don't know. We're going to do a whole show about it. He's at the Build Conference, which is where they're announcing this. I've downloaded it, and I have it installed and running in a virtual machine on my computer. But I think that this early version is a little buggy, this evaluation copy.

Steve: Oh, you think? We're going to have a field day.

Leo: Yeah. See, you and I are different kind of ends of the spectrum. I like the new thing, and you hate the new thing.

Steve: Well, I love iPads. I mean, for what it is, iOS is just fantastic. But the idea of trying to pad-ize the desktop that I'm going to be sitting in front of and want to get serious work done with, I just - no. That's just wrong. I mean, I hated XP. They just put sugar-coating on top of Windows 2000. But boy, what they did with Win7, I just - it just fights me. It's like...

Leo: Well, you probably won't like 8 because I think, as far as I can tell, it's Win7 with a tablet shell on top of it, which is probably from your point of view the worst of all things. But we should point out the reason that you don't like the new thing is not so much because it's the new thing, in general, it's because you want something tried, true, tested, bug-fixed, and secure; right?

Steve: Well, yes, there's a security aspect to it. But, for example, I run with an Explorer window open, showing me a tree view of my system. I mean, and I can instantly jump around. I've got everything organized. I know where everything is. The moment I start using Win7, I just immediately come to, okay, wait a minute, how do I get that? And it's gone. They took it away because they said, oh, people don't want that. Well, Steve Gibson does.

Leo: And who cares what Steve Gibson...

Steve: They didn't ask me.

Leo: Oh, Steve. Thank goodness.

Steve: The good news is that XP has three years to go, and I'll be with it to the dying moment.

Leo: By then they'll have figured out Windows 7, and you'll be set.

Steve: Yeah.

Leo: You skipped Vista. You're going to skip Vista, probably.

Steve: Probably? Oh, my god. Waiting was the best thing that ever happened. I can just completely jump over that nightmare. Yeah.

Leo: I wish - it would be kind of nice if they would, somebody would just kind of make a version that would be kind of like, for those of you who don't want updates, upgrades, improvements, this is the reference version of the operating system. You know, like DOS. DOS hasn't changed in 20 years. Be nice to have a version of Windows that just, this is it, we're going to leave it like this for those who want it. I guess that'll never happen.

Steve: Well, and you could almost do that because, I mean, the fact is, by the time that, like, three years from now XP will have settled down, and they will have stopped messing with it. Well, they have stopped messing with it. And they will have, however, been dealing with any backfilling, any old security problems because it'll still be under security support. And at that point all of the new stuff will be what's under attack. It's very much like people using Windows 98 2nd Edition, 98 SE. Boy, I mean, nothing infects that now because it's of a different entire DNA strain.

So it may very well be possible to just stay with XP. The problem is that developers will be developing for lord knows what we'll have then, Windows 26, and it won't even - none of their software will run on XP. And that's going to be a problem. People who want to stay way far behind with something that just works, that gets the job done, we won't be able to. We'll be forced forward because applications will be forcing, or new peripherals or whatever. I remember, for example, like remember USB was sort of forcing us forward because only the later Windows platforms 10 years ago supported USB. And if you wanted to have peripherals that were USB, you kind of had to move forward. So it's inevitable. But still, some of us will just be bringing up the rear, and that'll be me.

Leo: Proudly.

Steve: Happily

Leo: Proudly bringing up the rear, Mr. Steve Gibson, ladies and gentlemen.

Steve: Absolutely. Unabashedly, without apology. I know what's on my hard drive.

Leo: All right, Steve.

Steve: Okay. So I was pounding for several weeks on Apple, the fact that Apple had not taken the DigiNotar trust out of their root for Safari on the Mac, and anything else that was using trusted certificates on the Mac. And it was later that day that we recorded Security Now!, last Thursday because we swapped with Windows Weekly last week, that the announcement was made and the update was made available. It's 188K, so not very big. And it did exactly the right thing. It completely ripped all trust of DigiNotar out of the Mac.

I wanted to see it myself, to prove it to myself, so I went to DigiNotar's site. And I drilled down, followed a couple links, basically pretended to want to be buying an SSL certificate. That of course is the last thing I wanted to be buying from those guys, I would say those clowns. But it did allow me to get to an SSL page because of course they were trying to bring up their own security in order to give me an HTTPS connection to their server. What I found was two things. One was that the certificate itself had expired, interestingly, a few days before. And I thought, well, okay, that's odd.

But then in viewing the certificate details there was a chain of trust back up to the Certificate Authority. And I forgot to say that, before I applied this patch, I used Safari to do this. And I saw that I couldn't just bring up the SSL page because the certificate was expired. But if I looked at the chain of trust, the root CA, the DigiNotar root was still trusted. So then I thought, okay. So that's established. Then I applied Apple's patch and did the same thing again. And sure enough, as we would hope, not only was the certificate itself expired, but when I then examined the chain of trust, the root was no longer trusted by Safari. So that was my way of verifying that Apple did the right thing.

I tweeted the link that I found that afternoon because it's a way of allowing people to verify that their Macs and their Safaris no longer trust DigiNotar. Unfortunately, this was a little convoluted because the certificate itself was expired, and you had to look at the root to see whether that was not trusted. And it's hard to convey that over a 140-character Twitter tweet. So some people were confused. People were trying iOS or their iOS-based devices, their iPad or their iPhone, which have not been updated. A good friend of mine commented that as far as he understands, iOS cannot currently be incrementally altered, that is, all the updates that we've had of iOS have been "Here's your new iOS." It's, what, it's a 500MB blob that you're forced to download. That'll be one of the good things they're doing with iOS 5 is they'll finally be giving us incremental updates.

So it may just be that iOS is not going to be changed because Apple already plans to change it shortly, and they don't want to be forcing another half a gig download just to

remove a single certificate from the iOS store. So I didn't even think yet to try my iPad, but I haven't seen any updates that would have changed iOS. So as far as I know, it still trusts DigiNotar stuff. The good news is Mac and Safari don't after last Thursday's small 188K update.

And there has been confirmation from GlobalSign that they were breached. We'll remember that the hacker known as Comodo Hacker, whose lengthy Pastebin online post I read to our listeners last week, he was claiming to be in four other major Certificate Authorities, and he named GlobalSign as one. Upon hearing themselves named, they suspended certificate issuance and hired the same people that DigiNotar hired. However, GlobalSign has been a model of proper behavior for this sort of problem, well, saving the fact that they were apparently breached.

They've confirmed that one of their web servers, the web server that hosts its website, was breached. And they said, "As an additional precaution, we continue to monitor all activity to all services closely. The investigation and high threat approach to returning services to normal continues. All forensics are being shared with the authorities and other CAs to assist with their investigations into other potentially related attacks." GlobalSign said that "system components" were taken offline, but started to come back online with the help of Cyber Defense Institute Japan on Monday.

So they haven't said, which means, because I think they are acting responsibly, they do not know of or suspect that any certificates were issued behind their back. It looks like some sort of web server breach did occur, and they're sharing details with other CAs to help them spot the same sort of problem. But it doesn't look like we have a problem with GlobalSign, certainly not like we had with DigiNotar, but not even like we had with Comodo before. Remember Comodo had nine certs issued without their knowledge due to that hack, and the serial numbers of those certs were then blacklisted by the various browsers quickly. But it was felt that Comodo did the right thing in notifying everybody quickly.

Two things in our Miscellany category. I spent about an hour Monday morning being interviewed by KABC TV.

Leo: Really.

Steve: Yeah, their consumer guy, Ric Romero, was doing a story on Password Haystacks.

Leo: Hey, that's great.

Steve: Really cool.

Leo: Kind of surprises me, to be honest.

Steve: Yeah. Well, I think it's because, if understood correctly, and I hope I was able to make it clear to a standard TV-watching audience, remember that the conceptual realization I had when I was actually working on something different, I was working on an interactive tool to help people design secure passcodes, that's when it hit me that

what really mattered was length over complexity. As long as the password you had wouldn't be found in a dictionary, and that meant you forced someone who really wanted to crack your account to do a brute-force attack, then what you want is the largest possible haystack for them to be searching, which means a longer passcode, or password, passphrase, whatever you want to call it. So Time magazine

had a column a couple weeks before which mentioned coming to GRC and the Haystacks. And the producer, Alan Gitterman of KABC here in L.A., he read the page, found it over on GRC, thought, wow, we should tell our viewers.

Leo: Wow. That's just amazing.

Steve: So in the next day or two, they're not sure when it's going to air, I will tweet when it does, for those who have a more real-time connection to news from me. But certainly the week after there will be a link to it on the KABC website for anybody who is interested. So that was cool. And it did force me to clean. If you notice, you can sort of see, like, back there, there used to be a visible pile of papers on the desk.

Leo: I know. You got rid of the microphone and everything.

Steve: Yeah, I did some housecleaning. However...

Leo: It's great. It's very pretty. Did they do it via Skype?

Steve: No, they were here.

Leo: Oh, they were there. That's cool.

Steve: And the cameraman is a listener.

Leo: Awesome. Of course he is.

Steve: And so it was very cool when Ric and the cameraman were walking up to the front door. He said - because I had already met Ric downstairs, and he was waiting for the camera to show up, who rolled up in the big bright blue KABC Eyewitness News van with the antenna sticking out the top and everything. And so when Ric and he were approaching, he said, "Steve, our cameraman is a regular listener to Security Now!. He knows all about you." So that was really - that was cool, and it made it also all very comfortable and sort of - Ric I think knew that, okay, this isn't, I mean, he didn't know where I'd come from.

Leo: Yeah, Ric's kind of a legend among Farkistas. But I'm...

Steve: Well, Ric Romero, I certainly recognized the name immediately when the producer...

Leo: He did a - but Fark gave him a really hard time because in 2005 he did a report on the "new phenomenon of blogging."

Steve: Oh. Well, new for his viewers.

Leo: But new for him. But he's made up for it because I'll tell you what, Fark doesn't even know about Password Haystacks, probably. So that's good. He's getting ahead of the curve now.

Steve: So my cleaning efforts over the weekend in preparation for Monday morning's visit were slowed down.

Leo: Uh-oh.

Steve: Because I could not stop reading. I sat on the Starbucks patio, I said, okay, I'm going to read till 9:00 a.m. on Sunday morning. Then I said, okay, I'll read till 10:00. And then that deadline passed. And I said, okay, I'll read till 11:00. Then I'll read till 12:00. Then I'll read till 1:00. And I ended up finishing the book.

Leo: Isn't it exciting, you've discovered the immortal words of Nathaniel Hawthorne finally. No? Not Nathaniel Hawthorne?

Steve: Not Nathaniel Hawthorne. Someone mentioned to me a new author and a series that sounded interesting. And so I tracked it down, put it into Amazon. And Amazon had the paperbacks, but nothing on Audible, nothing on - well, Amazon wouldn't tell me about Audible. But nothing on...

Leo: No, actually they would because they own Audible.

Steve: Oh, okay.

Leo: So if there's an audio book, they'll often say we've got the audio version.

Steve: Now I don't remember whether there was. But it certainly wasn't Kindle, which is what I was looking for. So I thought, well, okay, I'll just get the first one of the series - of 12, Leo. So I'm, like, oh, could this be good.

Leo: Damn you, Steve Gibson.

Steve: So the first book arrived in dead tree form, pulp. And I start reading it late last week. And actually I was having dinner with Jenny, and I left it by mistake. I was distracted and left, I walked away from the restaurant where I had been there reading and waiting for her and a friend of hers to show up.

Leo: Oh, dear.

Steve: So I got to the car, ran back to the restaurant, and in the meantime they had seen that I had left it and had given it to Jenny, who had driven off in her own direction. So I sent email because I was all set up to spend the morning on the patio reading some more because it was already seeming really good. So I sent her email saying, uh, I need that book desperately because, even if I bought another copy, I couldn't get it till Saturday. So anyway, I Googled the name "Honor Harrington eBook." And it turns out it is in electronic format. And moreover, the first two are free. And it's in every format known to man - EPUB, MOBI, PDF, HTML, I mean, and it's - so Baen is the publisher. This first book, Honor Harrington, it's titled "On Basilisk Station," was written in '93, I think it is. Maybe '92.

Leo: Before people start looking for the author Honor Harrington, you're talking about she's the heroine of David Weber's books.

Steve: She is the protagonist, yes, of David Weber. Now, I wanted to give credit for whomever it was that turned me on to this. I mean, because Leo, oh, my god. Now, I have to characterize it a little bit because it's different than the other stuff, but fabulous. So this morning I put Honor Harrington into my Eudora full-everything search. And up started popping references in the Security Now! email folder...

Leo: They've been telling you about this all along, Steve.

Steve: April '07, Serge Beaumont. June '07, Steven Thompson. November '09, Matt Horton. Something must have happened in December of 2010 because John Weatherby, Gary Berg, Javier Gordo, and Denis Sherman, oh, and Jaime, all mentioned it, one using the subject "Looking for a (very) large sci-fi story?" Then...

Leo: Now, wait a minute. This can't be it. No, that's not it.

Steve: Last month Mike Lopez, Tim Lahey, John Whitlock, G. Wade Johnson, Neil Laubenthal, all mentioned this series, and one saying "The ONLY sci-fi recommendations you'll ever need." So, and then Robert this month, earlier, said "Good sci-fi series in the same vein as Lost Fleet." Okay, so, okay. The guy is a fantastic writer. And I'm seeing words like "perspicacity" or "perspicacious" or words I haven't seen since my college skills review class in high school. Beautifully put together sentences. It's hard sci-fi. It is absolutely the most unsexist series I've ever read. Honor Harrington is a young woman who does a great job in the first book. I couldn't put it down because I just wanted to see whether the reviews, which were all glowing, could be true. And they are.

Leo: Now, this is neat because it looks like "On Basilisk Station," which is the first, right, that you can read online. You don't even have to download it. The text is online.

Steve: Yes. Everything has been converted to Audible. So for those who are Audible listeners...

Leo: Now I'm happy.

Steve: ...all of this is on Audible. The books get bigger toward the end. I was looking at their sizes in various - after reading the first one, I bought - I downloaded the second one for free. And then I bought - they sell them there for $5. However, you can also poke around the 'Net and find them, with authorization, free. Because at various points some of the hardbacks were published with a CD containing the full text of the prior books. So they're not locked up. None of this is password or, what am I trying to say, protected.

Leo: DRM, yeah.

Steve: DRM'd, exactly. Anyway, so now, they're heavy on character development and sort of the political interplay and the interplay of the characters. So it's much less torpedoes firing and...

Leo: Good. I never liked that much.

Steve: Right. And what this brought me back in mind of was, back when I had a whole company of people, like 27 employees, they did something very cool for me one birthday, which is I received the writers kit for "Star Trek: Next Generation." One of them pretended to be my agent. And acting as my agent, she contacted Paramount and got the writers kit.

And the last thing I was going to do was, I mean, it would have been fun to write an episode of Star Trek: Next Generation. But I did sit down and read what writers were given. And one thing that stuck in my mind was it made it very clear, it said, yes, we have photon torpedoes and tractor beams and beaming and a warp core, you know, you never want to let the - you may have to eject it if things get bad and so forth. But it said "Star Trek: Next Generation" is not about those things. The stories are about people. They're about...

Leo: Shocking.

Steve: ...human drama set in that, placed in that setting. So writers, would-be writers, we caution you, we're not going to give your scripts a chance to hit the screen if you get yourself all wound up in dilithium crystal breakdown. We just, you know, Scotty has a problem with dilithium crystals every so often, and they may need to be recrystallized. We don't want to know what that means. We're more concerned about whether Scotty's

going to pull this off, and Kirk is going to sleep with a green alien. That's more interesting to us.

So this is like that. I am rapt with this author, with his style of writing, with his development. He's already setting up a fantastic scenario in Book No. 2. Now I'm not in a hurry because, if I am, nothing else will be getting done here at GRC, and I want to get the Off The Grid stuff finished. There's been a little delay in that because I just had to read the first book.

Leo: Well, at least you're honest, now.

Steve: Now I don't want to finish them because I don't want them to be over with. And I'm feeling guilty because the only reason I allowed myself to read the first one was I thought it was only available in paperback, and I've been reserving everything that's available electronically for my stair climber to force me to work out. So now, if I refuse to allow myself to read these unless I'm on the stair climber, I'm going to be in fantastic physical shape.

So anyway, the good news is they are around on the 'Net. If you put in "Honor Harrington eBook" - there's also Wikipedia entries, of course, with links to CDs. And you can find the text, if you don't want to pay. I happily paid $5 for the other nine of them. And the author is still at it. Wikipedia made mention of the fact that another book or two are coming out in 2012. So this is all alive and kicking. And, oh, boy. You can get into it for free with the first book and see if it's your cup of tea. Because, as you said, Leo, it's easier than Hamilton, a little less tech-y than Peter's stuff. But, oh, really, really nice writing. And I really like the way he handles characterization.

Leo: Darn you. Darn you, Steve Gibson. And if our podcasts are late coming out for the next three weeks, you'll know why.

Steve: Oh, it really is good stuff. So I give it...

Leo: I have to admit, the covers make me think it's kind of a little bit of a romance novel. But it's not; right?

Steve: Absolutely not. In fact, we learn in the beginning of the second book that Honor has never really had any interest in guys. And she doesn't like women; she actually says that. But just she hasn't run across a guy yet that she's going to. However, I did see something in a synopsis later on where...

Leo: I figure he's saving his powder for later. Yeah, oh, yeah.

Steve: Oh, and it's been called - it's been reminiscent of the Hornblower novels.

Leo: Really. Well, see, I was a big fan. I love those kinds of novels.

Steve: Well, and it's like that. In fact, the technology uses Warshawski sails to sail along gravity waves in hyperspace. And we have her looking out of her viewport at the frozen lightning appearance of the sails as she's sailing through hyperspace. And it's not overdone. I mean, it's just, oh, it's fantastic. So thank you, all of you, who for the last five years have been trying to tell me about her...

Leo: Finally, he paid attention.

Steve: ...and this author and series. And I wanted to pass the word on to all of our listeners because absolutely worthwhile stuff.

And just a short note from Sam in Texas, who says, "Wow." He says, "I applied SpinRite to a very old hard drive that had stopped working completely. It wouldn't boot into safe mode, or any other mode, for that matter. And even though the drive had many unrecoverable sectors and one recoverable sector after using SpinRite, the drive now works. I'm very happy that this product works like it said it would. I mainly use it for drive maintenance. Keep up the good work. Sam in Texas." So, Sam, thank you.

Leo: Yeah. All right, Steve, I have questions. I know you have answers or you wouldn't have given me these questions.

Steve: And we got just a bunch of great stuff this week. So let's plow in.

Leo: "Fire when ready, Gridley." Question 1 from Tudor Gazdac - what a great name - about "Watch it work" and "Get it done." Dear Steve, I've been watching for a while about how Off The Grid actually works. This was a couple of episodes ago, your great kind of non-electronic password system.

Steve: Right.

Leo: That's why it's Off The Grid. And I noticed something weird. When I selected the option "Watch it work" - so we should explain. When you go to Steve's website, GRC.com, what is it, GRC.com/offthegrid.htm.

Steve: Yup.

Leo: It'll generate a grid, and there's two ways to do it: quick, or you can see it actually happen. That's the "Watch it work" thing. He memorized the position of some characters as they were generated. But after it's finished generating the entire Latin Square, I observed a change in the distribution of characters I had memorized. The final result was totally different, a totally different Latin Square from the one that it was working on initially. For example, I memorized the position of three characters, OmG - capital "O," lowercase "m," uppercase "G" - as they appeared in the first column and the first row. When the Latin Square got generated, the characters moved from their positions, the result being a totally new Latin Square. Is

this a bug? But it also can't be such a big deal. But it's something I noticed, and I thought you should know. Congratulations for everything you're doing. Tudor Gazdac.

Steve: So, Tudor, you and many people mentioned that and asked. And so two things happened. First of all, I changed the way it works because it was confusing people, so that it no longer does that. But it was deliberate and not a bug. It was actually doing it in both modes, but in the "Get it done" mode you couldn't see it happen.

Leo: You could see it, yeah, yeah, yeah.

Steve: Here's the problem. The way I was generating Latin Squares, where I'm moving through each cell and selecting from among those that are still candidates, that actually results in a tiny bias from if all of the cells in the row were chosen at once. There's a famous logical problem. You probably will remember this, Leo, and I'm not sure I'm going to get it right. But it's like there was a game show where you had to choose between...

Leo: Oh, yeah, yeah, yeah.

Steve: ...what was behind Door No. 1, Door No. 2, and Door No. 3.

Leo: So it's counterintuitive, to be honest.

Steve: Yes. And so...

Leo: But your first pick can actually affect your second pick.

Steve: Yes.

Leo: Which doesn't seem right because, you know, if you flip a coin, what happens in the first flip doesn't have anything to do with the second flip.

Steve: Exactly. And so I think, was it Alex Trebek, or I don't know who it was, but somebody, there would be some fantastic car...

Leo: I think they did Let's Make a Deal, is what they used as an example.

Steve: Ah, okay.

Leo: Monty Hall.

Steve: Right, Monty Hall. And so you'd try to guess, like, where the prize was. It was behind one of these three doors. And you would make your guess. And then, given that you weren't right, they would open one of the other doors to show - or, no, they would open that door, yeah, they would open that door and show that it wasn't there where you had chosen. And you had the opportunity then to change - no, wait. You made your guess.

Leo: No, here's the deal.

Steve: Okay.

Leo: So the player picks Door No. 1. Now, they don't open Door No. 1. That's not how - if you've ever watched Let's Make a Deal, what they do is they open another door, let's say Door 3, to reveal a goat. And then they say, do you want to stick with Door No. 1, or would you like to change your pick to Door No. 2? And if you study game theory, you apparently understand, because I've never understood this, that you should always switch.

Steve: Yes. And that is...

Leo: It's to your advantage to switch. Now, it seems to me it doesn't make any sense because whether there's a goat behind 3 or not, why should you switch?

Steve: Exactly. And what's really freaky, Leo, is you can do this with coins.

Leo: You can prove this, yeah.

Steve: You can prove that it actually matters. Anyway, so that's kind of what was going on with the way the grids are being made because the characters at the end have less choices because the characters that preceded them on the line have already taken those choices up. So, and again, it's like counterintuitive. But some of the guys in the newsgroup, when I was working on all this, demonstrated it. I produced a whole bunch of 4x4 Latin Squares, and they counted them up. And sure enough, they were not equally distributed. There was a bias to this approach toward generating them.

Leo: How interesting. So the Monte Carlo problem, which is the other way you describe this, bit you.

Steve: Yes. And so what I decided to do was, once I had arrived at a completely finished Latin Square, I would randomly permute every row and every column, so that I get an extremely impossible-to-predict, incredibly high-entropy Latin Square. Then I have

another empty grid, and I randomly copy the columns into randomly chosen other columns, in order to make a new grid. And then I randomly copy rows from this new one back into the original one in a random sequence. So that permutes all the rows and all the columns to completely remove the fact that there was otherwise a detectable bias in the way that these grids were being generated, even though it really doesn't matter. But you know I'm a perfectionist, and I want this to be just nailed down the right way.

So that's what people were seeing was they would, because the "Watch it work" mode plunks the characters out slowly, you could memorize the beginning of the grid as it's working down further down below. Then suddenly I would perform this permutation that completely results in yet another grid. It is a variant of the first one, and it's necessary to go through that first process in order to have a chance to get to all possible grids. So this just is a short little hyperspace jump, essentially, to the final grid, which people were seeing. It no longer does that because it really didn't matter for people who wanted to see it work.

Leo: Awesome. Awesome. Question 2, John Benson from Twitter, @JBenson2, wonders - I love the Twitter questions because they're 140 characters. None of this "love the show"...

Steve: We're about to make up for it, though.

Leo: Oh, got a long one.

Steve: Yeah.

Leo: Some program logins ask for my Twitter or Facebook username and password. Seems risky. Could you comment on Security Now!?

Steve: Okay. So that's a really good question. And I'm seeing it more and more, too. There was somewhere I was looking the other day that offered me the option of logging in using my Twitter or Facebook account. Now, hopefully what John is talking about, certainly what the page I was looking at was talking about, was something we have done a show on called OAuth, or Open Authentication.

Leo: So you should know that in the early days of Twitter they did not require OAuth, but they now do. So anytime you give your Twitter credentials to anybody, you're not giving it to anybody, you're giving it to Twitter, which is then sending an OAuth token to that third-party site.

Steve: Well, okay.

Leo: Twitter does not allow this third-party login anymore.

Steve: The way it works, though, you still, you should absolutely never put your Twitter

or Facebook or other site into a page other than Twitter or Facebook.

Leo: Somebody could trick you into saying, oh, we'll do it for you. But they can't. So if they're asking you for that, they're stealing it.

Steve: Exactly. So John, and anybody else who's confused about this, when you say yes, I would like to use my Twitter or Facebook authentication to authenticate to this other third-party site, when you click one of those buttons you will be taken to Twitter or Facebook. There you're logging in with Twitter or Facebook, assuming that you weren't already logged in, like statically. And then you're brought back. And so what the OpenAuth authentication protocol does is behind the scenes it allows that site you were going to authenticate to via a site where you already have credentials established, it allows them to exchange things. But you, yourself, should only ever provide that authentication to the site that you are intending to authenticate to. And then you bounce back to the third-party site. And he'll say, oh, okay, fine. Facebook has said they know who you are. Twitter has said they know who you are. And now we have a token that represents who you are. And it is completely secure.

Leo: Yeah. And you should know that I think both Twitter and Facebook have deprecated any other form of authentication. You have to - so if somebody's saying, oh, we'll do it for you, that's a lie.

Steve: Oh, yeah, yeah. You never want to provide your credentials to other than this site where you are known by that credential.

Leo: I guess they could spoof Twitter. But you should be able to tell that it's a Twitter...

Steve: Oh, you ought to have an HTTPS connection. You ought to be able to verify the credentials of the page by looking at the SSL certificate and go through the whole thing, yes.

Leo: Now the long one, from Simon Bartholomeusz from Melbourne, Australia with a troubling and true, because you confirmed it, tale of widespread, sophisticated, attempted social engineering exploitation. [Bang] Ow. [Australian accent] Steve and Leo, I'm a longtime listener of the show and have been dying for an opportunity - I'm not going to do it the whole way that way - an opportunity to write in.

Steve: And actually, he actually wrote "cue Australian accent."

Leo: Oh, that's not you? He said that?

Steve: No, that was him.

Leo: [Australian accent] It finally happened on a Saturday evening as I was watching the football. And he wasn't - he was watching Australian rules, no doubt. Go Cats. While watching TV, the home telephone line rang. Suspicious that it would be one of the usual telemarketing parasites, I answered the phone anyway, looking forward to the opportunity to vent some frustration. I guess the Cats were losing. Sure enough, I was immediately presented with a dial tone that transferred me through to a young man with an Indian accent called "Carl."

[Indian accent] Carl explained that he worked for an organization that had been monitoring the security of my computer - Oh, yeah. You know what, this is happening all over because I've gotten calls on the radio show about this - and could confirm that it was infected with malicious software. Now, of course Simon's a listener, so he immediately said it's a little odd that he could figure out which one, given I have a bunch of computers behind a router. He knew the name of my girlfriend, who lives with me, presumably from the phone book, and stated that "I am here to help." Yeah, right.

Having been a listener of the show, the word "scam" immediately came to mind, but I continued to indulge my curiosity. Carl told me to boot up my computer. And I guess he did it, which is great. While it was booting, he explained that 98 percent of the files on your computer are corrupted, whatever that means. He also explained that he worked for a company based in North Sydney that had been involved in monitoring security software that had been installed in my computer before it was sold to me. Quite farcical, since I was booting up a Mac that had been reformatted umpteen times over.

We booted up the computer into Windows (I'm a dual booter), and I was immediately told to hit Windows-R to load up the Run box and type "eventvwr." That was quite funny. He directed me to the application log and explained that all the error and warning events were instances of Windows telling me my computer was truly infected. Bugger all.

Then back to the Run box. I was told to type "inf." INF, of course, stands for "infected files." That's what I was told. "Can you see a lot of .inf and .pnf files in the INF directory?" I was asked. "These are corrupted files." I was asked to double-click one of the PNF files. When Windows brought up a dialog saying it didn't know how to open the file, this was taken as further damning evidence my computer had indeed been taken over. Amazed that the Russian Mafia could so easily have compromised my system, I asked Carl what to do to fix this seemingly dire, utterly hopeless situation. Sure enough, there was a solution: "Go to the Run box and type www.ammyy.com."

I obviously knew better than to go to a website I knew nothing about in circumstances such as these. At this point I told Carl as much, giving him a polite piece of my mind, and I explained that I would immediately ring the police. And it gets better. I was transferred to his superior, another man with an Indian accent, who transferred me to his superior, an angrier man with an Indian accent, who explained to me that he was happy to patch the police in right now. I said, "All right, go ahead, mate." Then after being on hold for five minutes I decided to hang up and leave it at that. 30 seconds of Googling ammyy.com made it clear that the initial goal of the scam was to load up remote-control software on my PC. In fact, Steve, you checked this, and that is a free desktop access program.

Steve: Yes.

Leo: What frustrates me about this nonsense is I know plenty of otherwise intelligent people who would probably fall for this sort of scam, unknowingly handing over their bank account details and openly handing control of their computer across to join the ranks of some botnet army. If I have a question out of any of this, it is what can ordinary people do to stop these people? I don't take sufficient comfort in the fact that I can protect myself. Security Now! has plenty of devoted followers. There must be something we can all do to proactively cut the heads off the hydra. Well, you know what happened when that happened.

Anyway, on a much more positive note, I just want to say thank you to both of you for your efforts on the show. My reading list over the past 12 months - Peter F. Hamilton, Daniel Suarez, and Fatal System Error - has been purely Security Now! inspired. That's amazing. We've taken over his reading list. Moreover, the show has further inspired me to take a much broader interest in computer science, with many hours spent playing around with MASM and better understanding some of the most fascinating technology phenomena of our time. I might even load up on vitamin D supplements. Man, this guy's devoted.

If either of you are ever in Melbourne, let me now, and I'll happily take you to a footy game, my shout. Kind regards, Simon Bartholomeusz, Melbourne. Well, good news, I'm going to be in Melbourne in November, Simon, and I'm going to count on that footy game. Because, you know what, footy, Australian rules football, is by far the best sport ever invented by god or man. With the exception, perhaps, of Quidditch.

So, Steve, I have gotten this call on the radio show numerous times. So it's not just Australia. They do this worldwide. Probably there's a large group of out-of-work call center guys who they just say, hey, in between supporting Dell calls, why don't you make these?

Steve: Well, yeah. And if you do Google, as I did, ammyy, up comes all kinds of people saying, for example, how Google does little summaries, and I'm looking here, it says, "Help, a guy just rang our phone and asked my wife to turn on our PC. She hung up. It rang again, and a guy said he was calling as some error..." Another one, December 8, 2010, someone claiming to be from Microsoft phones you at home and tells you their logs are picking up an infection from your computer, blah blah blah. I mean, there's, like, people are falling for this and getting caught out by it.

Leo: Well, and I don't worry about our audience. I know they are smart enough. But I do worry about the radio show audience. And I suppose I should just make a blanket warning every single episode because those are the people they're going after; right?

Steve: Yeah, yeah. And, I mean, I didn't know you could type "eventvwr" and immediately bring up Windows Event Viewer, but it works.

Leo: Well, they know.

Steve: And so it's a very clever way of getting buy-in from pure social engineering attack, getting buy-in from a Windows user. And here's somebody who says, first of all, we're monitoring your security, we're going to prove it to you, and he brings up this log. And the Windows Event Viewer is always full of a bunch of random nonsense that isn't working right, or something didn't start up, or tried to boot, or who knows what. So, I mean, I can totally see that they would be bringing a person along and allow someone to believe that their system was in bad shape. And who knows what happens, I'm interested to know what happens if you do download the remote desktop stuff and they probably connect to your system. Somehow they're going to have to get money from you, so they probably drag you along and say, oh, well, yeah, now we're taking over your cursor, and poke around and do magic-looking incantations and then say, oh, well, now we need your credit card number, and we'll fix the problem for you.

Leo: I had a call on the radio show in May from Mike, who was in Edinburgh, Scotland. And he had exactly the same call, but he had the good sense to record it.

Steve: Ah.

Leo: So you can listen a little bit, if you're curious. It's on YouTube, if you search for "Windows Service Centre Phone Call Scam." [Playing YouTube file] "I understand that you have received a call from the Windows Service Center. And we are the associate of all the Windows operating systems all over the world. And the reason why we are calling you, because we have been receiving a lot of errors and warning reports from...." We received a lot of error reports from the computer you're using. It's the same script, I believe - "which has been malfunctioning with the Windows operating system and the [indiscernible] still going on. [Indiscernible], okay?" "All right."

So then this is exactly what you described. It's funny, I mean, Scotland, Australia, I think it's all over the place. I'm sorry it's not louder. But you can hear Mike talking to them. "[Indiscernible?]" "Mike: Sorry?" "Can you see all the [indiscernible]..." It's such an elaborate scam. I'll let you folks go and watch Mike's video. It's "Windows Service Centre Phone Call Scam." And apparently there are quite a few of these on YouTube. If you'd like to know what these sound like, maybe just send a link to friends and family. Mike did a great job of documenting this with video and everything. It's fascinating, isn't it.

Steve: Well, and it's been going on for about at least three quarters of a year because here's one that was posted on December 8th of 2010. I just can't understand how these people haven't been shut down. It's just mindboggling.

Leo: I'll tell you how they haven't been shut down. It's international. In the U.S., if you call, if you get a solicitor's call, by law you can say put me on your do-not-call list, and it would be a violation of federal law for them to call you back - unless they're in Bangalore. Then what do you do? And I think one of the reasons this works is because cheap international calling using voice-over-IP, they can afford to do this. Because you've got to remember that it costs them both the time of a person to make this call, the phone call itself. They must get some pretty good paybacks out of it, I would guess.

Steve: Yeah. Wow.

Leo: Wow is right.

Steve: Anyway, I wanted to bring it to our listeners' attention just because it was from the bizarro world, from our perspective. But Leo, I'll bet it works.

Leo: It's the real world. Here's Question 4, a worried listener who has asked for anonymity. He's in India. He says: Steve, I'm a listener in India and follow your podcast very regularly. Thanks for the information - by the way, I just want to apologize for the bad Australian and Indian accents of the previous letter. Didn't know we had any Indian listeners. Oh, boy.

Steve: There's a lot.

Leo: I know we do. I'm just kidding. Thanks for the information you share as I drive around for long hours. I have a query that relates to TCP/IP, and I would be grateful if my email ID and name are not disclosed.

Recently I was approached by the police, who traced my IP address for some harassing emails sent to an individual. I was pretty surprised. I was sure it was a mistake. How can a security-focused listener fall for hacks; right? Well, it turned out that the authorities had traced one of many emails to my IP, and other emails are being evaluated. The only outcome possible is that all emails end up being traced back to my broadband connection or to random people. I'm hoping it's the latter.

I had an ADSL 2Mbps line that plugs into an integrated modem and wireless router. The model is a Beetel router provide by the telecom company, but I believe my question is independent of the model being used. I was using WPA/TKIP for protection - that's good. Made sure that all services - Telnet, HTTP, TFTP - were only available to the LAN, not the WAN. That's good. I thought I was secure, and I didn't bother to change the default router admin password. Big, big mistake.

I started exploring my router in more detail. From an independent Internet connection, if I typed the IP address of my router in the browser, I was prompted for the username and password and was promptly taken to my admin page, which needless to say required only the default password [angry frustrated sounds]. I don't know what the definition is of WAN for my router. My modem has an option for showing the password as plaintext - sounds like a terrible modem.

Steve: I know.

Leo: So from the Internet my password was visible, along with the hidden SSID. To make things worse, my ADSL PPPoE username is my_phone_number@my_service_provider_name. So if somebody hacked into my router, he now has my SSID, my password, my phone number - of course that could be used to obtain my address through social engineering - even though my IP is

dynamic. So now I have reason to believe my router has been hacked. Now I do understand I'm vulnerable, but the hacker needs to get in close proximity of my house. He needs to be on the WiFi in order to use my router for malicious intentions. My question to you as a security expert is as follows. I don't know, does he need to use the WiFi? That may not be the case.

Steve: I don't think so.

Leo: Yeah, I think he could do that over the Internet. Assuming my machine was not compromised by altering the DNS or other means, whether a router can be configured such that a hacker is able to bounce his IP traffic off my router so that it looks like I am the sender of the email. I explored the DMZ, NAT, and Virtual Server protocols. They didn't seem to fill the bill. The router, which is a Broadcom router with a fancy marketing name - so it's Broadcom based, I guess - also has a Linux prompt and has very basic commands, one of which is IPtables, which is a Linux firewall.

I am really going bonkers while trying to narrow down whether the hacker came in close proximity to my house or did the bouncing of packets remotely. Is the latter possible, as it is more convenient? And, if so, what protocol should I be looking into my router? VPN is not supported in my router, which was a suggestion provided by one of my friends.

Please note that cyber investigations in India are pretty new and nave, and authorities believe an IP trace implicates the router owner, even if you've never met the person who has received the malicious emails. Spoofing, ability to hack WPA the way I describe above are really beyond comprehension for the police, and any insight could you provide with regards to IP routing that affects my everyday use is going to be greatly appreciated.

Please do keep up the good work - boy, I feel for this guy - and hope you bump up this question for the upcoming podcast. I'm looking forward to your views on this unfortunate episode. It's really nerve wracking, and it's making me angry. He has changed the default password.

Now, I guess I have a question. They don't really need to hack his router to spoof his IP address, do they? Can't you just spoof that?

Steve: Well, yeah. No, because email uses TCP, and you need to establish a connection before you're able to send traffic over TCP. So unlike UDP, like DNS queries, where it's a one-way transit, somehow you actually have to have - the three-way handshake that we discussed last week has to complete, and that does prevent IP addresses from being spoofed. Now, he said, first of all, he said his IP was dynamic. So although typically IPs are not changing often, my first thought is, who had the IP at the time that these emails were being sent? Because...

Leo: Oh. Well, that's a good point. It might not have been him.

Steve: Yes, exactly. I mean, that would be the first thing I would say. However, the fact

that he was able to log onto his router's administration page from outside on the Internet, knowing his current IP, is obviously a huge concern. It's, for example, what ShieldsUP! is designed - GRC's ShieldsUP! service is designed to detect that, and would. So if he were to go to ShieldsUP! and just run a scan, he would come up with a red box on port 80 saying that you've got an HTTP server running on this IP address which is accepting TCP connections from anybody. And so...

Leo: Oh, that's interesting. So that's one of the things you'd see in ShieldsUP! is that there's access externally. And it shows up as port 80.

Steve: Exactly, yes, because it would just be...

Leo: It's a web page.

Steve: Yeah, exactly. He said he went to a different web browser and typed in his IP address and got his page. So that means he's got a web server exposed to the Internet, and ShieldsUP! would make that very obvious that that was going on. So if any other listener wants to make sure that they're not in the same situation, GRC's ShieldsUP! service, it's been around forever...

Leo: Love it, by the way.

Steve: ...designed to let people know. Now, as to could a router bounce traffic or be involved somehow, the fact that it's got Linux in it means that it's potentially got the power to do that. But in order to bounce a connection, you actually have to have a proxy of some sort. You have to have the router able to accept a connection and then reissue another connection on your behalf from its IP. That way the connection to a remote SMTP server would be terminated at the router's IP. And it's unlikely that, I mean, that's not typical software in a router that's available over on the WAN side. So it's unlikely that that would be going on. And by the way, I'm looking at the traffic to GRC, and everyone who's listening to the podcast just started using ShieldsUP! at once.

Leo: Including me, by the way. My ident port is turned on. We're using an ASG. The only thing bad about having an ident port turned on is then that does identify that there's a router at that IP address.

Steve: Exactly, it demonstrates that it's not just a disconnected IP. And the only possible problem would be if there was a bug in the ident server, like a buffer overflow, that someone could then go, oh, let's see if there are any known problems that the ident server has, that kind of problem.

Leo: Otherwise we're fully stealthed. Of course we're using ASG, so I think we're pretty safe. I'm not too worried about that.

Steve: Yeah. And of course ident open is necessary for some connections to be

completed on other ports...

Leo: Often chats and things like that.

Steve: ...as we've talked about in the last couple weeks, yeah. So my best guess is, first of all, it's nice that he knows that he was exposed. Hopefully he's able to, I mean, he can use ShieldsUP! to verify that WAN-side access is shut down. Certainly that's better than changing his password. And, wow, I would worry, if you believe you've turned off WAN-side access, but it hasn't been turned off. That's a problem. So there's that.

And secondly, it's really unlikely that a proxy is or was running in your browser unless it's very sophisticated. But he said his doesn't support VPN connections. And I can't imagine a WAN-side proxying service that would be part of a router, which is what would be necessary in order to loop traffic through someone's router. Just spoofing wouldn't work. So I think it's more the most likely scenario is that somebody else previously had his IP. And I would just tell the police, hey, verify...

Leo: It's dynamic, dudes, yeah.

Steve: IPs, they do change, and somebody had it before him.

Leo: There's nothing more scary than a government authority with half knowledge of technology. That's just terrifying.

Steve: Yeah.

Leo: You have our support, sympathy, and - so it's not possible to spoof headers, then, in an email? At least as far as the servers go?

Steve: You can spoof headers, but you're going to finally send it to the SMTP server that will then deliver it. And that last header cannot be spoofed because it's putting the header on and then sending it on.

Leo: So the intermediate servers could be spoofed.

Steve: Yes.

Leo: But not the final, the ultimate server, or the first server. See, could the originating server be spoofed?

Steve: It's actually a chain of servers where the first header is - it's a chain of servers where headers are successively added to the list. But the one that really matters cannot be spoofed, only the previous headers.

Leo: Okay. We're assuming, of course, that the police understand that, which doesn't sound like they...

Steve: Oh, goodness, yes.

Leo: They're not even close. At least, hey, they looked at the headers. Ben in Reading, U.K. wants to implement Off The Grid for Android devices. I'd support that. Steve and Leo, I want to create an Android application for generating those Latin Squares you use in Off The Grid, which can be then used to generate a password, as you explained in Episode 315. The original idea was to use a master password to somehow seed the generation of the grid. But following discussions on Episode 316, I realized that this would significantly reduce the number of possible squares.

The revised decision on the design of the application is, when it first runs, a Latin Square will be generated and saved. The square will be both generated and saved locally on the device. When the user wants to access it, they can provide a master password to unlock the saved square. Then they enter a website address which will use the square automatically to generate the password. The user will have the option to select a starting square from the grid after entering their password and the website address. That's nice. I like that feature.

The way I see it is in order for someone to obtain your password for a site, they would have to, A, have your phone; B, have the master password to unlock the square; and selecting a square to start with. I am not considering an extra level of security because there's only 26x26 different starting locations to choose from. Still, that's a fairly large number, so it would take some time, perhaps enough for you to change your passwords because your phone was lost. I realize if someone does steal your phone and knows what they are doing, they might be able to somehow extract the Latin Square, removing the need to even enter a password. But I can't think of a way to overcome that. Does this seem like a secure approach? Anything I'm missing? Is this how you would implement this, Steve?

Steve: Okay. So here's what I would do. First of all, he's right that using a master password just doesn't have enough entropy in it for the Latin Square generation. So, and some people have said, hey, will we be able just to use our own password to create a Latin Square? And I'm not going to allow, I'm not going to provide for that option because I really want there to - part of the security of the system is that there are just so many, 10^428 or whatever that number was, a ridiculous number of possible Latin Squares. That's where it does get some of its security.

So we can't use a passphrase to directly seed the generation of a Latin Square. What I would do, and what I will suggest Ben does, is use a good pseudorandom number generator, and I've got one that I developed in JavaScript, and that's public domain, and I'm encouraging people to take it and use it. That's the ultra-high-entropy pseudorandom number generator that we talked about a week or two ago, which is posted on my website. If you go to offthegrid.htm, down at the bottom is a link to a page that shows that.

So that generates 1,536-bit-based pseudorandom numbers from a key of that size, 1,536 bits, which when it's base64 encoded is 256 characters. So that's a big, long - that 256-character ASCII blob is your master password, or actually the key which then generates

the squares. That you could encrypt using a passphrase. So that's arrived at, that 256-character key is arrived at pseudorandomly, and that's what is used to key the Latin Square generator. That, then, you symmetrically encrypt using a passphrase. So the beauty of that is that, if someone got your phone, they could put in a passphrase, and it would decrypt to some key that would generate some Latin Square, and not yours. So they would get a Latin Square, and they would have no way of knowing if it was the right one. Only when you put in the proper passphrase would it decrypt that 1,536 bits into the proper key to result in the proper Latin Square.

So it's actually sort of nice because nothing is on the phone that can be taken from you. Your passphrase decrypts the encrypted key into something. Any passphrase will decrypt it into something. It just will always be the wrong something unless it's the proper passphrase. But even the wrong something will happily generate a Latin Square. It'll just be the wrong Latin Square. And then they're really up the creek because then they have to start using the wrong Latin Square to try to guess your passwords, and those won't work, either. So it's kind of cool.

Leo: Good solution. So I'm looking forward to this app. Ship it.

Steve: Yup. And I meant to mention also, many people...

Leo: I bet there's others, yeah.

Steve: Many people have expressed an interest in turning this into an app where you put in the domain name, and it gives you out the password. So basically it automates the whole Latin Square path-following process. And I like that because you have that for convenience. It's not so much Off The Grid any longer, that is...

Leo: No, it isn't, yeah.

Steve: It's now in your phone. But I like the idea then of also having the same Latin Square printed out and folded up and tucked away in your desk drawer or stuck in your wallet so you always have that as a fallback, doing the same thing that you have automated with your smart phone.

Leo: I use a web-based solution called SuperGenPass that hashes - you use a master password, which never changes. And it hashes that with the domain you're at, GRC.com, to generate a unique password. I don't know, I mean, it's not completely clear how secure it is and so forth. But it's good enough. I don't use it on the bank and stuff like that. But it seems good enough for me, and it's automated, and it comes - because it's web based I can use it on any device, including my mobile phone.

Steve: Right.

Leo: Moving right along, Question 6 from another Australian. We get - I love it that we have listeners all over the world.

Steve: We sure do.

Leo: It's fantastic. India, England, Australia, and Alabama coming up next. But first, Richard Wilkinson in Sydney wonders why so many certificate stores. What's the point of so many certificate stores? He's talking about, I think, CAs, right, Certificate Authorities.

Steve: No, like storage spots for certificates in a computer.

Leo: Oh, okay. Why isn't there just one for the operating system, and the browser queries that? Actually, that's how Apple does it. What is the advantage of having separate stores for each browser? If there were just one operating system managed store, one update would do it for all of them. That's a good question.

Steve: He's actually right.

Leo: That's a legit question, although I presume that Google wants their own store, and Firefox wants its own store, and now Microsoft, et cetera, et cetera, et cetera.

Steve: That's exactly right. I think there are two motivations. One is that some, many of these browsers are deliberately and by design cross-platform. Firefox runs on Mac, PC, Linux, UNIX, Sun machines. It's multiple platform. So they have their code base, which includes their certificate store, and so they're managing them as a whole. So when you add a third-party browser, it's often bringing along its own certificate store. Now, as you mentioned last week, Leo, it's not necessarily the case because, for example, Safari on Windows uses the Windows certificate store, just as IE does. And I think that Opera does also. But many times a browser, a third-party browser will just bring its own along.

And so on one hand, part of it is that they're often open source and multiplatform. And I think the second part is that it's in the politics of browser vendors, third-party browser vendors who are doing their own browser, part of their whole philosophy is, well, we're doing one because we want control. We want to decide what add-ons operate and how they operate. We want to control the features. And so I think they proactively don't want to rely on the platform's certificate store and whatever policies it may have. For example, Firefox on Mac was secure against the DigiNotar problem several weeks before Safari on the Mac.

Leo: Precisely. Precisely. So Firefox doesn't want to wait until Apple fixes it. They want to fix it.

Steve: Right.

Leo: It gives them a selling point, too. Not that they're selling anything. Okay. Question 7, Ronald Stepp, Enterprise, Alabama. He wonders about what he calls an "interesting potential furball service," in other words, dropbox client-side encryption. Just picked up on this article about a new service called Bitcasa. By the way, I know exactly where he picked upon it because it was TechCrunch that wrote about it heavily. In fact, yeah, he gives us the link. And interestingly enough, of course, it's a TechCrunch investment. And they wrote so many very positive articles about it, each of which ended with, and by the way, we have an investment in Bitcasa. And it's like, come on, guys, stop that.

Anyway, it's promising seamless and integrated unlimited infinite cloud storage of everything on your hard drive, terabytes or petabytes even. They say they will encrypt it client-side - so it's PEE or PIE, Pre-Internet Encryption - but that one way they minimize costs is they can find out which files are duplicated across different users' machines and then only store one copy of it. What? That makes no sense. But okay.

The security concern, and what makes it interesting, plus lots of people are pointing it out in this article, is this: If you're encrypting the stuff client-side, how is it possible they securely eliminate duplicates by comparing files encrypted with different keys? A lot of good posts there in the discussion at the bottom of the page, TechCrunch.com slash slash blah blah blah, thought you might get a kick out of it. Well, obviously that's patent nonsense; right?

Steve: Well, it can be done.

Leo: Really.

Steve: And I wanted - so Ronald has his name on this question, but I wanted to acknowledge all of the people who have tweeted ever since TechCrunch did this, people saying, hey, Steve, how is this possible? How can it be possible for them to claim that they cannot decrypt the files?

Leo: Oh, I think Stride got it in our chatroom already. I'm curious to see if he's right. Go ahead.

Steve: I'll be surprised if it could have been explained so quickly because we're about to have a serious propellerhead discussion.

Leo: Really? Well, he said what if you generate a hash before encryption?

Steve: Yes, that's part of the solution.

Leo: Okay.

Steve: But then you still can't - how could you only store one copy? That's the hard part.

Leo: Oh, good point. Oh, you're right. You could know it's the same, but you couldn't - yeah, okay, good. So that's part one, okay.

Steve: So I went to their site to track this down. And they make a big point of saying, first of all, they have 20 patents. Well, we're about to bust one of them, or maybe more than one, because I'll tell everybody how we could do this. And I'm sorry if they got a patent on it...

Leo: Well, patents disclose how you do it, so they're protected.

Steve: Yeah. But again, the problem is patents are also supposed to be non-obvious to someone trained in the art.

Leo: Oh, and you figured it out, yeah, good point.

Steve: So if someone just asks me how we do this, I'll tell them how we do it.

Leo: It's obvious.

Steve: But they didn't ask me. And it may - some listeners may feel that they deserve a patent after they hear how it can be done.

Leo: Okay.

Steve: So they make a point of saying that they cannot respond to subpoenas, that this stuff is truly secure. Even if they were coerced to provide anything to any government agency, they can't. Which means they do not have the keys, period. They're also saying that they are eliminating duplicates, and that the reason they can offer for $10 a month truly absolutely infinite storage is that so many copies of files are duplicated, they're not having to store duplicates. So on one hand they're saying they don't have the key, which means that what they're storing is pseudorandom nonsense that is completely opaque to them. On the second hand, they're saying that they're doing duplicate elimination, so that multiple people having the same file are able to store, are able to access it even though it has a single copy.

Leo: Boy, this seems like something we should make a contest out of for next week, like how would you do this? Because I'll tell you, it's obvious to you, maybe, but I'm wracking my brain now. I mean, the hash makes sense, but how do you store a duplicate?

Steve: Okay. So here's one way it could work. And I just thought about this when I was

assembling the questions this morning. And I thought, okay, well, I could solve the problem. So all users have a private key and a public key. And the public key, being public, is part of their account information. So on the client, on the user's computer, it generates an asymmetric encryption key pair. One is kept private; one is public and is sent up to Bitcasa. So first user wants to store some files up there. So his files are blocked at certain sizes. We've talked about blocking, which - take a large file, break it down into one MB or four MB or whatever size big chunks you want. That, prior to encryption, is hashed. So everyone agrees - so that's sort of the easy part, is we're going to take the block and hash it to create a fingerprint of the pre-encrypted data, the plaintext data.

Now we generate a random 256-bit key which we'll use to encrypt that block. And that random key is - the user has it, that is, the user keeps this random key and sends the block and the pre-encrypted hash up to Bitcasa. And they say thank you very much, we're storing this, we don't know what it is, but we have its fingerprint. And remember that the fingerprint, the signature, the hash, tells them it uniquely identifies it from all the other possible ones. But they still don't - they have no idea what the content is. They don't have a filename. They don't have anything, just a blob. So this user has lost no security.

Now, somebody else, User B comes along, signs up for Bitcasa, downloads the client, and says I want to store this big file up there. His client breaks the file up into big blocks, makes a hash, and this - here now there are a couple ways it could work. But suppose the hash goes first. Bitcasa has all these hashes indexed and says, hey, we already have a hash that matches. So we're holding a block which was encrypted under a key we don't have. But we know who does have the key. So they take User B's public key and send it to User A, asking User A to please encrypt the randomly chosen symmetric key for the following block under User B's public key, which nobody but User B can decrypt using their private key. And so then Bitcasa sends that back to User B, who decrypts this private key that User A had which was encrypted with User B's public key, decrypts it with his private key. Now he has the key to the block. He never needed to uploaded it. He merely needed to say I have the following hash. Now he has a key to the block, which gives him access to that file. So the block only needed to be sent to Bitcasa once, and then only tiny keys are sent around. And the problem of User A maybe not being online is solved by the fact that, in fact, very quickly, hundreds and thousands of users will all be sharing the same blocks.

Leo: It also, and somebody in the chatroom is pointing this out, and I think this is an interesting - tell me if this is not right. If authorities got any one user's key, they'd have access to that file from all the users.

Steve: Yes. Except that's not a liability because it's just like any other file system. So I'm trying to think. So it is the case that there's metadata which Bitcasa is storing which knows which users are sharing which blobs. But Bitcasa themselves cannot decrypt the blobs.

Leo: No, they'd have to get - so authorities would have to go to one of the key holders.

Steve: Yes, one of the key holders. And then the authorities would be able to, just like going to your computer - and we assume also that there's some login process and so

forth. So if that user didn't give up that information, then the authorities would still not be able to have anything. So it is possible to share hashes of plaintext and to arrange so that the randomly chosen keys are never in Bitcasa's possession, yet they could arrange a socially networked key-sharing service using public key technology that would make the whole system work.

Leo: So the scenario, for instance, that would be potentially risky for you as a user, let's say "The Hurt Locker," which is a film that the owners of which have been very aggressive about suing people.

Steve: Ah, very good point.

Leo: So I have a copy of "The Hurt Locker." I upload it, and 50 other people also upload their copy of "The Hurt Locker." The authorities can't go to Bitcasa and say, who are these people? But they could get a list of everybody. They could find - if any one of us says "Here's the key," they can then verify that all 50 people have "The Hurt Locker."

Steve: Right.

Leo: And from Bitcasa they can subpoena the information that these 50 people have this file.

Steve: Right. Now, using the scenario I painted, all I was trying to do was to solve what little we know about this. It's not even in beta yet. Hopefully they will...

Leo: Oh, here's a good one. Web7088 says, what if the authorities upload a copy of "The Hurt Locker"? That's actually a great scenario because now they have a key.

Steve: Okay. They have a key, except they would still need Bitcasa to tell them everybody who is sharing that block, and that they are one of the people sharing the block.

Leo: But Bitcasa does know that, so they could be subpoenaed for that information.

Steve: I don't see how Bitcasa could not know that.

Leo: So we don't know, we're guessing about how they're doing it. But assuming that they're doing this with this eminently reasonable way, if what your concern is, is that you want to keep a copy of a pirated movie that is a problem on your site, all the authorities have to do is upload a duplicate, and then they have a key.

Steve: Everybody sharing that file would be potentially taggable.

Leo: And all they then do is they subpoena Bitcasa and say, by the way, who else has this file?

Steve: Right.

Leo: Okay. Interesting. We don't know if that's how they're doing it.

Steve: Maybe it's even fancier, and I hope that - but all I was trying to do was to solve those features that seemed contradictory. And in fact we can see how public...

Leo: There is a way to do it, yeah.

Steve: Yes.

Leo: That's interesting. Somebody in the chatroom says, oh, well, then it's obvious this is being funded by the Recording Industry Association of America, the MPAA. That's how they can do it for free.

Moving along, Question 8. Tim Miller and his son, in the Bay Area of California, rely upon, from last week, possibly hackable Medtronic insulin pumps: Steve, I'm a SpinRite owner and user, listener since the beginning, et cetera, et cetera, et cetera. Love the podcast. I'm a Type 1 diabetic on the MiniMed pump. Regarding Medtronic's defensive reply about users hearing the pump's beep, while it's audible, I don't always hear it. We were talking about the fact that you could hack the pump remotely. Not too remotely, but you could hack the pump. And Medtronic said, well, yeah, but we have audible beeps whenever the pump is reprogrammed, so the user would go, oh, yeah, somebody's messing with my pump.

While it's audible, I don't always hear it. My 10-year-old son is also on the MiniMed. He ignores or doesn't hear it beeping half the time either. My wife is the one who says, uh, I hear beeping, you should check your pumps. It beeps for all warning signals, by the way, which include low battery, low insulin, reminder to check your blood sugar, et cetera. So of course I'm concerned about this possible threat. The serial number that's required for the threat is a 10-digit alphanumeric number. However, it appears that four of the digits are always alphabetic, maybe the model number since mine and my son's are the same, and the other six characters appear to be numeric. So if the hacker only has six digits, and if the pump replies, this is a very fast crack to get into a pump.

Leo is correct about the pump range. You do have to be close to the insulin pump for radio reception. I've found the max to be about 20 to 30 feet. If an attacker were to shut it down, then any diabetic that tests regularly every two to six hours would notice their blood sugar rising without the insulin provided by the pump. But if an attacker could induce the pump to deliver a large dose of insulin, that of course could be fatal. And the way the pumps beep, you would typically not know until it beeped when done delivering the fatal dose of insulin, or when your blood sugar had dropped dangerously low. This situation could be life threatening.

Just thought you should have a few more details. And I will be contacting Medtronic to get a better response than, "Oh, we're not worried about it." Boy, I'd be worried about it, yeah. I mean, you'd have to, somebody would have to be wanting to get you. But it's just one more, you know, I can see there'll be a movie of the week based on this one.

Steve: Exactly.

Leo: From Hyderabad, India, our second writer from India, Gopi Krishna Reddy Guntuku wonders how a Certificate Authority can revoke a certificate that's already been issued: Steve, I love your show. I've been following from Episode 1. I understood almost everything that you explained about DigiNotar except for one thing. How can a CA revoke an already issued SSL certificate, technically? The CA signed the cert, which is given to the website owner; right? When an end-user connects to the web server, handshaking will occur. And during this process the end-user's browser isn't going to go to the CA website to check if it's revoked. But if that's the case, the browser can check if the cert is valid or not by this call, as well. Why would you trust to store it locally if you're not going to check? And what if the CA site has been hacked, or the end-user's browsers got attacked by DNS spoofing? Am I thinking impractically? Please clarify. That's a good question.

Steve: It's a great question. And we've touched on it a little bit here and there. We'll be talking in great depth next week, returning to the whole issue of the PKI, the Public Key Infrastructure, where we have defined the roles of Certificate Authorities and chains of trust and so forth, because there have been some proposals recently for ways of addressing this increasingly creaky foundation which relies on every single Certificate Authority in the world that our browsers are trusting all being completely perfect in their behavior. And we've seen very recently that's not only difficult, apparently it's impossible.

But part of this model has always been the notion of revocation. The need, the recognition was always there that a certificate might need to be revoked. The person whom it was issued to might be breaching the agreement under which the authority signed it. They might be misbehaving. They might be conducting themselves poorly. They might have lost control of the certificate themselves so that somebody else has it, and they might say, please revoke the certificate. We were hacked. Somebody may have gotten ours. We'd like another one. But in the meantime, definitely kill off the one we were using before because we don't want anyone being able to masquerade as us.

So there's many scenarios where there might be some need for not an entire Certificate Authority to be untrusted, as we have just seen with DigiNotar, but for specific certs to be untrusted. Now, we saw this with Comodo where nine certificates were issued fraudulently. And in fact it is because there really isn't a really well-functioning solution to this that the serial numbers of those certificates are now embedded in the source code of Firefox and Chrome and other browsers.

There are two solutions that are part of the actual structure, but unfortunately they don't work reliably enough. One is called a CRL, a Certificate Revocation List, the idea being that certificate issuers also maintain a site, a URL, typically under their domain, where by automation any browser is able to pick up a list of the certificates they have previously issued which have not yet expired, because remember that's where the expiration of

certificates comes in handy. Every couple of years we who use them are forced to jump through hoops to renew them. But the good news is that keeps these lists of problem certificates from having to grow forever because they only have to revoke those that haven't revoked themselves, essentially, by virtue of having expired.

So the CRL is an example. And just when I was preparing this question I put into Google "VeriSign Certificate Revocation List." And if you Google that, it'll immediately take you to VeriSign's very nice page where they list all the different types of certificates they issue and the URLs of their CRLs, their Certificate Revocation Lists for each type of certificate.

And in fact, for example, I was in Firefox. I clicked on one of the links, and up popped a dialogue with Firefox saying, oh, I'm adding this CRL to my store. So Firefox was happy to add that computer-formatted CRL, that Certificate Revocation List, to its knowledge of certificates that were revoked, although that was a manual process. One of the fields in certificates is a link to the signing authority's revocation list. So it is possible for a browser to be told, prior to trusting any connection, check for revocation. So that if I were to go to somewhere that wanted an SSL-secured connection, the browser could be configured to look at the certificate which the server has provided, asserting its trust. And in that certificate will be a URL which the browser could immediately then visit in order to see whether, by serial number, the certificate it is in the process of considering trusting is, as far as the Certificate Authority knows, still trustworthy. Have they declared it revoked or not?

The problem is that's a slow process. That requires additional steps every single time. Now, you can also cache this knowledge, so you would trust for a while, which of course does create a window of opportunity where there could be some exploitation. That's one of the two approaches. The second is something called OCSP, which is a certificate revocation service. And under Firefox, if you're curious, you can look, Firefox can be configured to make an OCSP query on a per-certificate basis.

The advantage there is that there's a third party maintaining revocation lists, and the browser makes a query by serial number, saying is this serial number revoked. And then the OCSP protocol will say no, it's not. Or there's even an option that Firefox has for not trusting certificates that aren't affirmatively confirmed, that is, if the OCSP server that you've configured in your browser doesn't respond, normally browsers fail open, meaning fail in a trusty direction, rather than in a not-trusty direction. You can change that behavior in Firefox.

So there are mechanisms, complex and messy. They slow things down. But no one has ever figured out any other way around it because we have this infrastructure right now where we're sort of statically trusting authorities that are in our certificate authority store, and we're trusting, as long as nothing has expired, the things that they have asserted we should trust. And the problem is, since that signing is done, things can happen to cause certificates to be revoked. The problem is incrementally, like, verifying every single time we make a connection introduces so much additional overhead, browsers don't typically do it. The systems are there in place, and hyper-security-conscious people could tell their browsers I really want you to be really secure, but it would slow everything down.

Leo: Wow. I think that, you know, it's interesting because this is what the Comodo hacker asserted. I have broken SSL. Is no good. Certificate no good.

Steve: Well, he certainly has given us all something to think about.

Leo: Yeah. Maybe not so good as we thought. Last question, actually correction, from Jim Schimpf of Derry, PA: I've been listening with interest to your development of the Portable Sound Blaster, aka the Dog Killer. The development board you found is a huge value. The Xpresso board; right?

Steve: Yup, the LPCXpresso.

Leo: For only a little bit of cash. I'll have to look into this for some of my projects. But the Arduino is definitely not in that machine's class. It's a 16MHz, 8-bit processor. But I did want to let you know that it is not interpreted code. You do have the full Atmel GCC tool chain on an Arduino, and you are compiling C or C++ code. What makes the Atmel chip an Arduino is a boot ROM and a loader. It's that plus the boot ROM and loader. You build code that's compiled and linked. If you use the tool chain outside of the Arduino IDE you can compile and link multiple files, then loaded via the boot loader.

Steve: Yeah. I wanted to make the correction. A number of other Arduino users were more knowledgeable about it and said, hey, Steve, you mentioned that it was an interpreted platform. And they're of course correct. I chose the ARM Cortex M-3 because it runs at 120MHz and is a 32-bit processor and, for example, has a single-cycle 32x32 multiply. So it's an incredibly powerful little chip. And amazing that it's $29.95 for the full development platform and all the software.

I poked around to understand what it was, how I got the impression that Arduino was an interpreted platform, and it's just that they use the term "sketches." They talk about "Arduino sketches" as being these, essentially the source code. And they do have a wrapper which sort of attempts to make the platform easier to program. Instead of where in a normal C program you have a main subroutine, and main is invoked by the loader at runtime, and then it's pretty much up to you, well, it is totally up to, typically, what happens, in this sketch mode with the Arduino platform you've got two routines, a setup and then something called "loop."

And so you put in the setup subroutine anything that you want to happen first during initialization at the beginning. And then the loop subroutine is just continuously called. And so you put in there all of the code that you want to sort of keep alive and have operating and then - but within that it is just a regular compiled C/C++ framework. So not with the overhead that I was assuming of interpretation, just a much less powerful based hardware platform, but also certainly there's a lot of support for that, as well. So I just wanted to make sure everyone had the clarification.

And the guys at the LPCXpresso site, they got 50 in, sold them out, got 50 in, sold them out, got 50 in again, and they were, when I looked this morning, about half sold out of that one. So I'm just tickled that we've got so many listeners that are interested in poking around with hardware because it's just - it's a world of fun.

Leo: Yeah, I think that you've started something there, and I think that's kind of cool. We want to do a show, by the way, kind of about this stuff. Not a robotics show or an Arduino show, but just a maker show, where we talk about people doing things like the Portable Dog Killer and robots and 3D printers and all of that stuff. I think it

would really be cool. We want to build a space in the back here. We have a little lab area where we can do that, be kind of fun.

Steve: Cool.

Leo: Well, Steve, that's it on the 10 questions from our great listeners. Boy, you all are fantastic. We appreciate those questions. We do this Q&A show every other week. So if you want to get questions in for two weeks hence, all you have to do is go to GRC.com/feedback, and Steve's got a form for you. Much easier for him to handle than email. GRC.com/feedback. GRC, of course, is the spot where you can get 16KB versions of this show, transcriptions, the show notes. Steve's done a great job there.

We also offer it at TWiT.tv, and you can watch live on TWiT.tv every Wednesday at 11:00 a.m. Pacific, 2:00 p.m. Eastern, 1800 UTC. If you're at GRC.com, don't forget to check out SpinRite, the world's best hard drive maintenance utility. It's a must-have, if you've got a hard drive. GRC, that's Steve's site. You can follow Steve on the Twitter, he's also there, @SGgrc. And you've been using Twitter a lot. Now we've got to move you to Google Plus.

Steve: Yeah, I'm there, but I just haven't really figured out what it is yet, so...

Leo: It's just longer.

Steve: I'm busy...

Leo: And more conversation.

Steve: Yes. I'm busy reading David Weber, our newly discovered author, thank everybody again for that. Wow, I mean, great, great space warfare stuff, space opera.

Leo: So I'm wondering if I should get an Audible version of it, look at the free version, get it on the Kindle. I can't decide.

Steve: It's all there for you, Leo. I mean, you're an Audible user. Just try the first book. And I tell you, you've got - you're opening a serious, beautiful adventure. And, oh, I should say this author does something I was worried about. I mean, it's why I really wanted to vet him first before I mentioned him, even though it was that afternoon that I began reading him, after the last podcast, so I didn't really have any opportunity to mention him before now. But it annoys me when authors of what is clearly a series spend a lot of time in the follow-on books telling us all about everything that we already know from the previous books. It's like, yeah, okay, I know, I read that one. And I read the one before. I don't want to - and this guy, actually I've been very impressed because I'm sensitive to that. He just beautifully sort of in context tells you only what you need to know.

So if you've picked it up in the middle - and why would anyone do that when all the books are there? Start at number one. But if you did for some reason, only Book 5 washed up on the island that you were stranded on, then you beautifully get only what's necessary. Sorry, you didn't read one through four, you're out of luck on that. But he does fold in enough just beautifully. So it's not frustrating. It's in context. It's only what you need to know. I'm just - I'm going to have to stop reading this on my Kindle. I'm going to force myself back onto the stair climber so that I'm going to...

Leo: What a good excuse to get some exercise.

Steve: Oh, god. And I think I'm about two thirds of the way through "FreedomTM." And again, I'm only allowing myself to read that when I'm working out. So now I have my next series; and, oh, it's going to get me in good shape.

Leo: Steve, you're the greatest. Thank you for being here. I've got some reading to do. So do you.

Steve: And next week we're going to talk about, yup, next week we're going to talk about how we come up with alternatives to what we've talked about several times just now in this podcast, the whole certificate authority infrastructure.

Leo: Oh, good.

Steve: How do we solve these problems?

Leo: Steve's modest proposal.

Steve: No, no, not mine, other people's. Because everybody knows we've got a big problem here. So I'll be doing research on that and present it to our listeners next week, Leo.

Leo: There's got to be a better way. Thank you, Steve.

Steve: Thank you.

Leo: Thank you all for joining us. We'll see you next time on Security Now!.