Listener Feedback #181

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-439.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-439-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We'll talk about the NSA speech by President Obama, and Steve will give his report card. We'll also answer some of your questions, too. It's a jam-packed Security Now!, up next. Stay tuned.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 439, recorded January 21st, 2014: Your questions, Steve's answers, #181.

It's time for Security Now!, the show that covers your security and privacy online with my good friend, Explainer in Chief himself, he's waving and pointing at his head, Steve Gibson.

Steve Gibson: I am a talking head. There's nothing below here. I'm mounted on a tripod.

Leo: Do you wear - I never asked you this. Do you wear shorts when you do the show, or do you wear pants at all?

Steve: Yes, in fact, I'll never forget the time I stood up to check something or get something, and you said, oh, you have no pants on. And I said, Leo, on the audio podcast that really won't - that won't come across very well. So better, better...

Leo: Very funny.

Steve: Yeah, I have shorts mostly because I get worked up during the podcast, and I'll get too overheated otherwise.

Leo: Really.

Steve: So we kind of have, yeah, we kind of have a - yeah. We have kind of a gloomy day today, so it's staying cool. It's 74 degrees in the office, and some doves are building a nest...

Leo: [Cooing]

Steve: ...outside. Yeah, they're kind of nice to hear in the background.

Leo: So today's a Q&A. We haven't done one in ages, it seems like.

Steve: It is. It was funny because I had to go digging back because we actually did a raft of them toward the end of the year. We were doing them to sort of catch up. And then we had all of our New Year's things going on, and then we had our New Year's catch-up episode. So, yeah, #181 for Episode 439. And we'll talk briefly because there's really not much to say this week about Barack's Friday announcement, what I call the "do nothing" NSA speech. We've got some news on the famous Target retail PoS malware that we were expecting. I want to revisit CryptoLocker a little bit because some numbers have come in on just how much of a windfall those bad guys made. I'll talk a little bit about my TSA interview that I did last Wednesday and about Evan's note for a superior alternative that means you don't have to go to Sacramento. And I have verified you can just go to SFO.

Leo: Yeah. Evan sent that note to me, too. And, yeah, I'm curious.

Steve: Yep. Also a little, sort of in miscellanea, I found a fabulous free email archiving solution for those email packrats among us. I'm one. Eudora, which I'm still using, just began to collapse under the weight of, I'm not kidding you, a quarter million pieces of email.

Leo: Well, and Qualcomm stopped updating it; right? I mean, it's orphan software.

Steve: Yeah, but it works fine.

Leo: I guess email doesn't change that much.

Steve: Just like my XP. And also some news about SQRL's password encryption technology. I've become an expert on the Scrypt algorithm, which we talked about in a podcast many moons ago [SN-388]. Remember memory hard algorithms, which resist

being accelerated by GPUs, FPGAs, and ASICs chips. And of course it's a Q&A. So we've got, actually, nine. I don't know why, but nine felt like the right number.

Leo: Given all the other stuff, you may be lucky to get to that.

Steve: Yeah, so a great podcast.

Leo: Very good. I'm excited. So, let's see, Steve. We should probably talk about the news.

Steve: Yeah. I spoke many times last week about the upcoming speech or presentation or whatever it was, pronouncement, that we were expecting from Barack Obama, the current U.S. President, obviously, about his administration's reaction to the NSA. We were hoping for something. And essentially we got nothing. The EFF, I suspected, would have a nice analysis of it, and indeed they did. I created a bit.ly shortcut for people who are interested because - and you might want to bring this chart up, Leo: bit.ly/ all lowercase b-o, as in Barack Obama, hyphen n-s-a [bit.ly/bo-nsa]. And this is their scorecard of his speech. Basically, they gave him - they had 12 points, 12 major issues that they were hoping he would make some movement on. And so they could have all been ones, in which case he would have scored a 12 out of 12. As it was, at least half of them were zeroes, and there were only a couple that were ones. And actually they were both, I think both of the ones were FISA comments. One was oppose the FISA Improvements Act, which he is doing, and reform the FISA court, which he is doing. Those were the two ones.

Overall he got a 3.5 out of 12, not even out of 10, out of 12. So, and as I was listening to it, I was in real-time during the speech, most of it seemed just like punting, like he was saying, well, we will be in the future doing this, and we will be in the future doing that. We're going to have a committee do this. And a lot of stuff was being turned over to Congress. So it's like, well, okay. I guess, you know, I'm sure he did what he felt he should. But mostly it was just sort of a nothing speech. So I wanted to follow up on my having mentioned that that would be happening on Friday. And it's like, yeah, okay. Maybe he did all that he could. I don't know.

Leo: I have to think that you get in the - in fact, somebody said this, you know, because he was very, as a Senator, aggressively outspoken against this kind of government surveillance. And somebody said, imagine day one, you just got inaugurated, you come, you sit down in the Oval Office, and they say, "Mr. President, here's the daily threat assessment." And you go, oh, my god. We've got to do something about this. So I hate to rush to judgment, given that we don't, we just don't know how much risk there is out there. But he has paid lip service, and I think it should be more than lip service, to the notion that we've got to balance protection with freedoms and constitutional freedoms. What do you think about this, the NSA doesn't get to - somebody's going to collect the data, but the NSA doesn't get to store it. They have to petition for a request. Does that not seem like an improvement?

Steve: It does. And really the way I think this should be set up, the only way that it is feasible is if there is a burden imposed on the providers. And they don't want...

Leo: They don't want this at all.

Steve: Yes. They don't want it. But I say, tough. Look at what you're getting. I mean, the amount of revenue you're generating and the piss-poor way you're treating your customers, typically. I mean, the idea that I'm not using all my bandwidth in a month and so it resets at the beginning of the next month, and the idea that in many places I'm being charged for text messaging, I mean, that's just ridiculous. I mean, anyway, they've got a sweet deal.

So given the cost of mass storage, the fact that cloud storage is essentially free for individuals now because it's so cheap to purchase drives, in return for the privilege of having the sweet deal that the providers have, they archive these records. And they don't have to go back forever. Ten years, for example, is fine. And it is metadata. It is not large data. It is tiny little entries of a start and an end of a call and what account it was and what cell towers it was on, basically the metadata. For one thing, it'll compress like crazy. And tell them they have to store it for 10 years. That's their obligation, in today's new world, if they want the privilege of being able to provide this service. And then they have to have an infrastructure that allows them to respond in obviously automated fashion to approve requests for a query from that data. That's the only way I can see that there is a believable wall between the government wanting this kind of data and the people who are holding it. If a third party is set up, that will just be seen as a...

Leo: Yeah, let's let Target store it. Why not, yeah.

Steve: Precisely. Well, or...

Leo: It's got to be the phone companies. It's got to be the telcos.

Steve: Well, or it'll just be seen as a surrogate for the NSA storage. It's got to be fragmented. The NSA has to put out a request to all the providers, and then they provide what they have, and then the NSA links it together. And it's worth mentioning, also, that it's standing out that, for all of the furor that this metadata collection has generated, there has been no clear terrorism win in return for all of this. That is, it turns out to be one of the weakest sources of information in terms of actually delivering verifiable results from what we know.

Leo: Well, but again I acknowledge that we probably don't know everything, and maybe they can't talk about the wins for a variety of reasons. I mean, one of the things that, you know, you can't - this has always been a problem with espionage and crypto, in fact. Remember with the Enigma machine it was a real issue for us, the Allies, to use information they got from the German Enigma machine during World War II because it would be, in effect, an admission we'd cracked it.

Steve: Right. They had to, like, set up fake convoys that would just by chance encounter these ships at sea and then radio in, in order to explain how they had the information.

Leo: So, I mean, look. I don't want to be an apologist, but I also understand completely that nobody wants to be the guy who said we're going to stop all of this spying, and then there's a terrorist attack the following week. And then of course the howls for saying that you've hobbled the agencies, and you're not letting them do their job, and of course we got attacked. This is very, very difficult.

Steve: I think the whole thing is a foreseeable, reasonable set of outcomes. That is, again, as we mentioned last week, Edward Snowden is hated within the NSA. I would say of course he is. He is heralded in other freedom-loving organizations. And again, of course he is. I think that what we saw was an overreach in our own kneejerk reaction post-9/11. It was please don't ever - you can have anything you want so that this never happens again. And then we now understand what that means better. And so there will be some compromise. There will be some negotiation where we find a middle ground. And I also think ultimately we're going to get used to this. I think we're going to be surveilled.

Leo: It's the new normal.

Steve: Yes.

Leo: Well, and I think that's the bigger question, is given these technologies, the Internet and telecommunications, the kinds of technologies we rely on and we use every day now, is our conception of privacy really kind of antiquated? Because, if you're going to use these technologies, it just kind of goes along for the ride that privacy might not be possible.

Steve: Back in the analog cell phone days, I remember that I had my Jeep Cherokee with the cell phone in the center console where you picked it up, and it had a spiral cord coming off of it because all the equipment was, like, stored under the back seat, and it had the antenna, I mean, the old-style analog cell phone. And I remember also having a scanner at the time, which scanned those cell phone frequencies. And it was moderately interesting to listen to some of these conversations. These were conversations that most people didn't understand were radio. I mean, it was in the clear. It was right there.

And my point is that many times I would be having a cell phone conversation with my attorney, and I would say, "Okay, I will call you back when I get to the office because I'm heading there now, and I don't want to talk to you about this now," because it was just radio. And so back then we were using a technology that was convenient, and it was incredibly insecure, so that anybody just scanning could be listening to one half, you only got one half of the conversation, which was sort of interesting because you could sort of guess what was going on on the other side of the conversation from the part that you could hear.

So step forward now 20 years, and we have the Internet and essentially the same sort of thing. If you want to have a non-private meeting, a multi-way meeting over the Internet, you certainly can do that, with the understanding that it may not be as private as you think. If you absolutely want privacy, then the five of you go meet physically somewhere and have a private conversation in a room that is hopefully secure. So, yeah, I mean, it's a different sort of paradigm. And it is the case that we've lost this notion of this being a

private medium.

Leo: Yeah. It's very challenging. And I think we speak a lot about - and certainly people who do this show and all of our audience, we're libertarians, and we don't want government intervention. But I think it's important to raise the point that these are not, I don't think, evil people. They're doing what they believe is right, and they are patriots, I would say almost certainly. And the issue is that we do have a Constitution, and we do have some rights, and we've got to fight for those rights. And I agree with that, as well. So it's just - it's tough. And I think that part of the problem really may be this larger global thing that we've moved into a world of tech. This is a change, as you said, this is inevitable given what's happened in technology, that, A, because you could collect all this and analyze this massive amount of big data, that somebody would. And it was just inevitable. And so maybe we have antiquated notions of privacy because it's crazy to assume that you can use all this stuff freely and get all the benefits of it and not be spied on.

Steve: Well, and this, too, is no big surprise. It's always the case that the social side lags behind the technology side. Technology is zooming ahead, creating new capabilities. And it's the understanding of it and the social adapting that takes time. I think that kids who are growing up now who have never known a pre-Internet world just sort of assume that stuff is being monitored. That's just - it comes with the territory. So their whole life, as they're living it over the next hundred years, will be framed differently than ours was, where we actually used to imagine that paper mail coming to our mailbox hadn't been opened and inspected, and that there wasn't a keyword search being done on all of our mail as it went to and from for the purpose of monetizing our eyeballs with relevant ads connected to the email that we were reading. But that's just - that's today's world now.

Leo: What a world, what a world.

Steve: So we do have some more information about the point-of-sale software, malware, which was installed in the 40,000 Target point-of-sale machines, apparently. It was originally generated by a 23-year-old Russian youngster. And apparently the guy who sort of has been initially targeted was some guy named Sergey Taraspov. And he's, I think, 17. And apparently he was doing, like, tech support for the 23 year old. And this 23 year old has formally confessed to being the original developer of this. The Target point-of-sale devices are - wait for it - XP.

Leo: Of course they are.

Steve: They're running a modified version of Windows XP, the embedded version, which essentially is a sort of a toolkit. I remember looking at XP Embedded, like, years ago, wondering if maybe that would be a useful platform for SpinRite to run on, if I could get a license for that. But I'm not paying Microsoft any per-instance licensing, so, no, that didn't make sense. But I did learn about it. And it's essentially - it allows you to produce an XP from, like, a componentized model, where you only choose the components of an XP OS that you actually need in your embedded application. So it ends up being much smaller and arguably has a much smaller attack surface. Unfortunately, it still does run XP applications.

And so since it's crossing over into XP's world, this is a piece of malware which runs on XP, which is able to do what's called "memory scraping." Essentially it opens the PoS process, thus gaining access to the point-of-sale process memory. And it just does a - it does memory scraping. It searches memory for the credentials while they're briefly in RAM after the card has been swiped and the user has entered their PIN code. It captures those and pulls them over into its own process space. So we know that this thing is called BlackPOS. And we've got our weekly trash pickup happening next door.

Leo: We're here to pick up the trash. You just, you know, anything you want to throw in here, bills, letters, anything, documents, passwords, you should go right ahead and throw those in there.

Steve: So as late as January 16th no antimalware software is recognizing this. So it's been known for, like, it was first seen back in September of last year, of 2013. And it is believed that Target discovered it around mid-December because it was briefly uploaded to the Google-owned VirusTotal site and appeared there and then was taken down not long after. So there's a research lab, Seculert, that found the sample and actually executed it under a virtual environment of Windows XP. And they discovered that it has two stages. It first infected their checkout counters, their point-of-sale devices, to extract credit card numbers, and it collected them for six days. Then it uploaded those to another machine in Target's network. And I did notice some reports saying that part of the way they got in was poor passwords. Apparently the internal passwords were easily guessable, and so the software used those in order to move the collected customer data onto the central server after six days.

And then it was exfiltrated to another website somewhere else in the world, and that location was never given. And that appeared to be a hijacked website that was running an open FTP server. So that FTP server collected this data. And then a third virtual private server located in Russia was used to download that stolen data from that hijacked intermediate server over the course of two weeks, pulling a total of 11GB of stolen sensitive customer information over the course of that time. And these guys say that there was no indication, given the FTP logs, the only connections they saw were to Target servers, or from Target servers to this FTP server. And they didn't see any indication that the also-suspected Neiman Marcus compromise was also going on at the same time. However, in very up-to-the-minute rumors which are beginning to surface, there are apparently six other retailers that have been identified, but not disclosed, that also appear to be victims of this software. So that's happening.

Leo: You want me to mention the garbage man again? If you're hearing - it sounds like perhaps Steve has something going on behind the scenes. It's just the garbage guys. They'll be gone in a minute. Yeah. There's always something. When you live in suburbia, you've got...

Steve: That's right.

Leo: ...your garbage men. You've got your lawn...

Steve: We had construction work going on this morning. And I thought, oh, lord, how long is this going to go on.

Leo: It's fine. It's not...

Steve: Yeah.

Leo: It's fine. Don't worry about it.

Steve: It's a real environment.

Leo: Exactly.

Steve: So, CryptoLocker. I wanted to touch back on CryptoLocker because Dell SecureWorks, they have something that they call the SecureWorks, I mean the Dell that we all know about, something called the Counter Threat Unit Threat Intelligence. And what I liked about this was they had a very sort of executive summary they put together for discussing CryptoLocker. So they said, as far as the crypto side goes - and I liked this because it was very clear, very succinct, and also used the latest information that we have.

They said: "Instead of using a custom cryptographic implementation like many other malware families, CryptoLocker uses strong third-party certified cryptography offered by Microsoft's CryptoAPI. By using a sound implementation and following best practices, the malware authors have created a robust program that is difficult to circumvent. The malware uses the Microsoft Enhanced RSA and AES Cryptographic Provider to create keys and to encrypt data with the RSA and AES algorithms," as we talked before, public key and private key technology.

"The encryption process begins after CryptoLocker has established its presence on the system and successfully located, connected to, and communicated with an attacker-controlled command-and-control server. This communication provides the malware unit with the threat actors' RSA public key, which is used throughout the encryption process."

Many people have had questions about which drives CryptoLocker would infect. And this made it very clear: "The malware begins the encryption process by using the GetLogicalDrives() API call to enumerate the disks on the system that have been assigned a drive letter, e.g., C:. In early CryptoLocker samples, the GetDriveType() API call then determines if the drives are local fixed disks or network drives, either DRIVE_FIXED or DRIVE_REMOTE, respectively. Only those two drive types are selected for file encryption in early samples. Samples since late September also select removable drives, which can include USB thumb drives and external hard disks, as well. After selecting a list of disks to attack, the malware lists all files on those disks that match the 72 filename extension patterns that the drive encrypts. Over time, the threat actors adjusted which types of files are selected for encryption. For example, PDF files were not encrypted in very early samples, but were added in mid-September."

Okay. Under "Action": "Each file is encrypted with a unique AES key." And that's obtained from Microsoft's cryptographic provider random number call. So it's unfortunately a high-quality random AES key. And that key is then in turn encrypted with the RSA public key received from the command-and-control server. Consequently, due to the nature of public key crypto, you have to have the private key to decrypt that AES key. And that is

never on the computer until the person pays the ransom.

"The encrypted key, plus a small amount of metadata and the encrypted file contents, are then written back to the disk, replacing the original file." And as we said a couple weeks ago, there's like a little header shim added to the top of the file. And then also - oh, and then it explains that: "As a form of bookkeeping, the malware stores the location" - and I haven't seen this anywhere else, by the way, so this is new information - "the malware stores the location of every encrypted file in the Files subkey of the HKEY CURRENT USER\SOFTWARE\CryptoLocker registry key." And that may also be CryptoLocker_0388. But so there's a registry key there that has the path of every file that was encrypted. And apparently that's what the decryption software enumerates in order to go find all the files that it encrypted in order to decrypt them after you've paid your ransom.

Then, after finishing the file encryption process, and this was also important: "CryptoLocker periodically rescans the system for new drives and files to encrypt." So if anything comes along, you plug in a USB drive, if the drive is there during one of these rescans, it'll grab it and encrypt those files, as well. "The malware does not reveal its presence to the victim until all targeted files have been encrypted. The victim is presented with a splash screen containing instructions and an ominous countdown timer," and so forth. And finally it talks about payment options, which we have only had fragmentary coverage of, but we'll wrap this up with an amazing analysis of the amount of bitcoins and their current value which have been transacted, because that has been tracked back.

So, payment: The ransom amount varied in very early samples and, as we know, we covered this, settled at what was essentially $300 U.S., or two bitcoins, back in the early days when CryptoLocker was introduced and a bitcoin was valued at $150. But in their analysis they say: "Dramatic Bitcoin price inflation in the later months of 2013," which we've talked about, "prompted the threat actors to reduce the ransom," first to one bitcoin, then to half a bitcoin, and then finally to 0.3 bitcoin, where it remains as of the date of this publication. And that's where it is now. So, and then this talks about how various payment options were offered, and ultimately MoneyPak and Bitcoin are where things settled out. And apparently the malware does a scan for the CryptoLocker command-and-control server every 50 minutes, waiting to verify that the payment has been accepted. And then, when it has, it'll then obtain the private key and start the decryption process.

So the reason I've been saying since this first arose that this was unfortunately going to set a precedent for the future is just how profitable this has been. Ziff-Davis did some research using some of the bitcoin chain tracking software, tracing four addresses which were used and were determined from multiple CryptoLocker victims who, after paying their money, made public the address that they had sent payment to. The CryptoLocker extortionware acquired a total of 41, just shy of 42,000 bitcoins, 41,928 bitcoins.

Leo: Wow. Wow.

Steve: Yes. And, see, that makes sense when you multiply it by, from another standpoint, the known number of infections was somewhere between 200,000 and 250,000. So 41,928 bitcoins at today's value of $960 per bitcoin, means that this generated more than $40 million.

Leo: Wow. A million dollars a day, says Violet Blue on...

Steve: 40 million. Yes, in fact in one day there was - they tracked - in one day they tracked a million dollars' worth of bitcoin payment.

Leo: All going to the same bitcoin address.

Steve: Yup.

Leo: Wow.

Steve: So the bad news is all these guys did, remember, this was written in C++, so they used the built-in API, the Cryptographic API sitting there in Windows. Didn't have to bring any crypto themselves. I mean, that's trivial to do, but they didn't have to. It's all there in Windows. They simply wrote some software that did the crypto right, and they used the state-of-the-art botnet technology we've talked about where, based on the calendar, a large set of candidate domain names, random-based, again, cryptographically derived, date-based, a huge variety of date-based domain names, one or two of which were actually valid, making it very difficult for bad guys to track them down. And so basically a state-of-the-art network arranged for command and control of this. And they netted themselves $40 million. What this means is we've not seen the last of it. And this unfortunately is - it's too lucrative not to be copied.

Leo: It's not - it's interesting that that was a single bitcoin address, implying that that was a particular, one particular person. But there's no reason to assume it's just one person doing CryptoLocker. Or is there?

Steve: Yeah, I think it's one person.

Leo: Just one guy.

Steve: Yeah, well, or one organization. I mean, it was well written. It may be organized crime in Russia. But it's believed to have Russian roots from looking at the code. But I think it had a single origin. There was a copy that we talked about, not written in C++, written in Delphi I think it was. And there are some others which there's some buzz in the online forums of other things coming, but that haven't happened yet. And there's been some not-quite-done-as-well me-toos already. But it just makes too much money. This is way more profitable than putting some clickware on someone's machine that's going to click on ad links. This is the big-time now.

Leo: All the major antiviruses recognize it now, and of course awareness is raised. So I presume it's slowing down. I don't know. Is it?

Steve: Oh, yeah, yeah. Yes, in fact, there are some charts that show the rate of infection, and it is now way, way down.

Leo: So you make your money quickly and get out.

Steve: Yup, exactly. Well, I mean, it's, look, many months. And it was getting, it was biting a lot of people. So there will be another one, and it will not be seen by antimalware, and it will just - all it has to do is copy this. Do the crypto right. There's been total coverage of it. So even someone who has no idea how to do it right, now they know how to do it right. It's just, as I keep saying, it's just not hard to do this anymore. All of the technology is available.

Leo: Yeah. I love that they're using a Microsoft library for the crypto.

Steve: Yup, yup. And I ran across this in some of this coverage. Carbonite, one of this show's sponsors, was reported in November to have been dealing with several thousands of phone calls from CryptoLocker-infected victims because, remember, Carbonite does not map a drive letter, which means that the Carbonite data stored was not accessible to CryptoLocker. And so what you wanted to do was you wanted to make sure that Carbonite didn't back up your encrypted files for you; or, if it did, that you went back to a further, pre-encryption version. But now Carbonite has a dedicated team for dealing with CryptoLocker recoveries, meaning that they have trained...

Leo: I hate to say it, but it's good for their business.

Steve: Yeah. They've trained up a subset of their tech support people to specifically help people with recovering from CryptoLocker.

Leo: Wow.

Steve: Yeah. So we have a friend of the show, you and I, Leo, Evan Katz, who I think he hails from the East Coast, but he was out here in sunny L.A. when he sent you a note and copied me on it. He travels a lot. And so he said: "Re TSA pre-check," he said, "you probably do not want to get it," essentially addressing you, Leo. He said: "Rather, you and your family and friends should do the global entry program run by the Department of Homeland Security," which is of course distinct from the TSA. Evan continued, he said: "The reason why you want to do global entry and not pre-check is a global entry includes very expedited coming back into the U.S. from overseas, and pre-check does not." And so, you know, I'm not a big overseas traveler. But so certainly, when you've been leaving the country, Leo, that certainly makes sense for getting back.

He said: "Moreover, the DHS program has many more locations that you can go to which will be much closer to where you live." And I verified that. We talked last week about how your closest TSA PRE was Sacramento, whereas I verified that DHS has an office at San Francisco Airport. So way more convenient. And he says: "And the DHS program costs only $15 more than pre-check." That's correct. Pre-check is $85; the DHS program is $100 for the same five years. Anyway, so he says - he signs off saying hope all is well

with you and wish you all the best and so forth.

So I did, myself, I did my TSA PRE check screening Wednesday, so I'm now officially - I get something within 21 days in the mail. And they fingerprinted me. Interestingly, my left hand wouldn't fingerprint at all. We tried, like, eight times to get my four fingers of my left hand to register, and they wouldn't. And what I noticed was on the screen it was doing a feature extraction in real time, right there. So, for example, first I did the four fingers of my right hand. And it took a couple tries, but then it worked. And then they wanted both thumbs at the same time, so I gave them both thumbs, and that worked. But when they wanted the four fingers from my left hand, and I'm left-handed, maybe I just don't have any fingerprints over there. I don't know. And I was looking at the screen. You could see, like, it imaging. And there was just nothing recognizable as fingerprints coming up on the screen, no matter how, I mean, he had me try eight times. And then he said, okay, you don't need them. I was like, oh, okay, fine.

So it was sort of strange. And so basically it was an interview. It confirmed the information that I filled in online. And I'm annoyed that this only lasts five years because I'm probably only going to travel five times in the next five years. I go home for the holidays, basically. But still, it's such a fabulous program, especially for anybody whose business has them traveling a lot. I just can't recommend it highly enough. It is an absolute win.

Leo: They want to know what - so where did you do your Q&A for the TSA PRE?

Steve: I had to go up to Long Beach.

Leo: It's not an airport. You go to a government building.

Steve: Yeah, oh, and, boy, this thing was really low rent. I was thinking, okay...

Leo: No, you want government buildings to be low rent. Understand, we're paying the rent.

Steve: I am glad I'm there during the day and not at night. It was Suite E105, and there was like a tractor-trailer dumping ground next door. It was in the - and there was, like, old oil wells being pumped on the other side in Long Beach. It was way - it was the wrong side of the tracks.

Leo: And what did they ask? What kind of questions did they ask? How do you feel about al Qaeda or anything? No.

Steve: No, nothing, not at all like that. They asked me to read the things that I had already filled out online, which were where were you born, you are a U.S. citizen, I mean, nothing about my philosophy, nothing about my past. Basically, are you a citizen? Is this your current address? And that's it. I mean, it was just sort of to look at me in the flesh. There was a camera aimed at me, so I assume at one point either I was being videoed or a picture was snapped, as is the case with a passport. And they were - it was

like a fingerprinting center. They were doing this sort of thing, identity screening, for other organizations or agencies also. So it was like it was a multipurpose setup.

Leo: So presumably the real check is against you and no-fly lists and criminal record checks and things like that. That's all done in an automated fashion.

Steve: Well, and what's so bizarre, too, is that people are reporting that they're getting random TSA PRE on their boarding passes. Just sort of...

Leo: Shhhhh. Shhhhh.

Steve: Yeah, like when Jenny said she never went through it, it's right. She just lucked out and got random PRE on her boarding pass.

Leo: [Whispering] That never happens.

Steve: Oh, okay, yeah. And my point is, as a security guy who's always been screaming about how ridiculous it is that I have to take my shoes off and be patted down and be body scanned when I'm flying 500 miles north to San Francisco, here they're just sort of saying, oh, you know, don't bother with that. We've selected you at random from a list, and you don't have to do that today.

Leo: Weird.

Steve: It's like, what? Then why does everybody else have to? Oh, it just makes my blood boil, the whole nonsense of all this. Okay. So, Leo, you're going to need - center yourself securely over your ball.

Leo: I'm on my ball. Okay, yes?

Steve: My copy of Eudora, which I'm still using on XP, was beginning to have problems with the - it might have been the 53,000 pieces of Security Now! email.

Leo: You don't even do email, I thought.

Steve: No, no, that's where the Q&As come from.

Leo: Oh, the Q&As go there, okay.

Steve: Yeah. So I have received 53,000 pieces of email from Security Now! listeners over the course of this, what is it, 9.5 years or something. Or maybe we're ninth year,

and it's 8.5 years. Anyway. And in addition to that, I had 200,000 other pieces of email because I don't ever want to throw them away. And it's very handy to, like, what was that something or other? And so I go searching through it and find something. Anyway, Eudora was just collapsing completely under the weight of this and had been getting worse and worse and worse for maybe about the last year. Finally this weekend I thought, okay, I have to do something about this.

My point is, the point for bringing this up is I am now, oh, my god, I am in heaven. I found something free which I can recommend to all. I know among our listeners there are people like me. There are people who just want to be able to find a piece of email that they're sure they received or sent 10 years ago, and whatever means they have of managing it now might be causing problems. So MailStore.com, m-a-i-l-s-t-o-r-e dotcom.

There's two versions of this thing. There's Mail Store Server, which is their commercial side, which, eh, it's a couple hundred dollars for five client licenses. And then there's Mail Store Home, which is completely free. And it is unbroken completely free. I mean, it works great. It's what I've ended up using because it was enough for me. I didn't like Mail Store the commercial version, only because it occupied, like, 50 to 100MB for the server component always running in the background, and then the client when you ran it. And I didn't want this thing running on my server because I like getting my mail off of the server, just for security's sake, and having it all be on my own workstation. So basically this is a very nice...

Leo: This is a good idea. I like this.

Steve: Yes, a very nice indexer. So either the home - the commercial version has some additional features. The home version will pull from Microsoft Exchange, Google Mail - now, there's an example, Google Mail, all these people who are using Google Mail, and Google has all your mail. Well, all you have to do is turn this thing loose on your Google Mail account. It will suck it all out of Google and index it and store it for you on your own drive. You can tell it where you want it to go, so it could be on an offline drive or some other drive than your normal work drive or on a network drive. Works with IMAP and POP3 and others. For local systems it can pull from Outlook, Outlook Express, Windows Live Mail, Thunderbird, SeaMonkey, and also EML, MSG, and MBOX files. So Eudora stores everything in MBOX files. So I was able to just have it - basically it just sucked everything out of all of my various folders in Eudora and indexed them very quickly. You can also export, without any encumbrances, anything that it has archived. So you can put stuff back out of it to Exchange mailbox, IMAP, or mail it to any email address via SMTP. And it also supports Outlook, Outlook Express, Thunderbird, and SeaMonkey for exporting in addition to these various email file formats.

Anyway, I love it. So what I have now is Eudora is stripped back down and is running at full speed again because, for example, all of those Security Now! emails are indexed. And the indexing system and the searching system is spectacularly fast. I can put in any - oh, and it understands regular expressions and so forth. I can put in any phrase that I think I dimly remember was in an email somewhere, and almost instantly I am looking back through time at those things where it occurs. And I know that contemporary email systems offer that feature, too, and certainly you can search your Google Mail and so forth. But if you'd like the control of a very nice indexing/archiving system, I can't recommend this thing highly enough. And it's free. So I just wanted to share it with our listeners.

Leo: I'm looking at a variety of similar programs. A lot of them are commercial. A lot of them do cloud, which of course you wouldn't want. But the reason I am is because I'm on a Mac, and this is Windows only.

Steve: Ah, okay.

Leo: And there are some open source solutions that are cross-platform and so forth. I mean, the idea is fairly simple. You pull down all the mail and store it in a file and then index it. But so I think this is a great idea. I found one called Got Your Back that's an open source Gmail backup program that works cross-platform. But I'm sure there's others, as well. So, good. A great idea. And, you know, I just leave everything in Gmail because that was the promise of Gmail originally is just store everything there.

Steve: Now, where is - everyone knows your email address, Leo. That's not a secret, I realize. Where is all that email?

Leo: Well, that's what I'm saying. It's on Gmail. You just leave it.

Steve: Even though you don't have a Gmail address?

Leo: Oh, yes, I see what you're saying. So the address I use actually goes through two servers. It goes first to - where does it go first? I think first to Gmail for spam removal, and then to another IMAP company, a commercial IMAP.

Steve: So you had Gmail pulling that account.

Leo: It pushes to Gmail. So basically you hit Leoville.com, the server just says, you know, it has an MX record, says I don't do mail, but these guys over here at Google seem to. And so everything goes boom, bounces off my server to Gmail, where it's stored. And then I pull it from Gmail. I can't remember which - truthfully, I don't remember who gets it first. But both have it. And Gmail has such good antispam features that I always run it through Gmail. But that way you can keep it forever on Gmail. But you keep it on Gmail. Be nice to have a local copy, I think. Yeah.

Steve: So Jenny and I saw a movie yesterday. We're in miscellanea time, if you hadn't realized. I just wanted to give a shout-out to the new Tom Clancy movie, "Jack Ryan: Shadow Recruit." I loved it. I mean, and I also loved "Patriot Games" and "Hunt for Red October" and "Clear and Present Danger" and "The Sum of All Fears." So consider that. If you are a listener who also loved all of those, and that kind of movie, you've got another treat in store because this was just, oh, my goodness, it was rip-roaring wonderful. It was just spectacular. And, boy, Keira Knightley is easy to look at, too. So, yeah. Recommendation.

Leo: Good.

Steve: GRC is about to get true quantum random numbers.

Leo: Ooh.

Steve: Yes. This came a couple days ago. It's from a company in Finland that produces a USB dongle - and here is a picture of it, holding it up in front of the camera there - which uses, we've talked about this before, it uses a reverse-biased semiconductor junction to generate wide-band Gaussian white noise. The noise is amplified and digitized through an A/D converter. The raw output bits from the A/D are then further processed through an embedded microprocessor to combine the entropy from multiple samples into each final output bit, resulting in a random bitstream that's practically free from bias and correlation. And so what I'm going to do is that will be plugged into GRC's server, and the server will be pulling from it. The data rate from these is typically not super high. If you want really super high, then you can spend more money. But this is several hundred thousand pseudorandom - or, I'm sorry, I'm in the habit of saying "pseudo." This is not. This is the holy grail of random. This is absolutely quantum phenomenon unpredictability.

So what GRC's server is going to do is fill up a big ring buffer of this. And then anytime someone goes to GRC's random, you know, the Perfect Passwords page, I will pull a set of truly random numbers from this generator and use them to seed the existing algorithm, which is now just running forward. And so essentially every single person who goes to the page will have truly random numbers that seed this otherwise really good pseudorandom number generator because I want to be able to handle a high traffic level, which that page is now generating. But the pages will, as soon as I get this installed, and this will be some time in the future, true random numbers. And I will let everybody know when we switch over. Actually, I'll put a picture of this on the page and change the diagram to show that we're actually pulling from a stream of the universe's entropy.

Leo: Isn't that a beautiful thing.

Steve: So I got a nice little note about SpinRite. And the guy, for whatever reason, said "Please keep my name off the air for this testimonial." And he just said, "Thanks again for a great tool. I use it on all my drives personal. It has even brought back some drives that were getting recycled. If I came across a disk that failed a DBAN DoD wipe, I would run the disk on SpinRite at Level 1 until completed. Then, if it failed DBAN again, I would run the disk through SpinRite at Level 2. If again the disk failed at DBAN, it would be run at SpinRite Level 4. Only 1 drive out of 100," and he says "literally, needed a drilled-platter treatment. For every other dead drive, SpinRite was able to bring it back to life for DBAN wiping."

So this guy had drives that had died that he had compromising information or private information, I mean, and just really any drive that dies you probably don't want anyone else getting access to that data. So it was often the case, since the drive had died, that DBAN - and by the way, DBAN is Darik's Boot and Nuke, D-a-r-i-k-'-s, Darik's Boot and Nuke, which is as it sounds, a bootable media. I think you can set it up for either CD or USB. And it boots a little environment that runs this program that will do a relatively good wiping of your drive. The problem was these drives had died, and so DBAN wouldn't

run. So he wasn't using SpinRite to bring them back to life to use them, but mostly to be able to run DBAN on them.

And so the good news is, not long from now, that will no longer be necessary because there will be a product, an inexpensive product from me, which has already been named, and I've had the trademark for it for years, and that's called "Beyond Recall." And so what the plan is, as soon as I get SpinRite 6.1 finished and out the door, I'm then going to basically take the core of that new technology that I developed to make 6.1 run so fast and repurpose it as a GRC-grade drive-wiping tool, which just ought to blow the doors off DBAN and everything else in terms of its performance because it will use, for example, the 32MB buffer technology that I've got running already in the work that's been done on SpinRite 6.1, where we can do multiple terabytes in the course of a few hours. And so it will bring that kind of wiping, and actually it'll be the second commercial product that GRC has.

Leo: After all this time, that's only your second product?

Steve: Yup.

Leo: Isn't that great.

Steve: It's funny, I was considering and proposed putting it into SpinRite. And the chorus was unanimous among all of the guys hanging out in the spinrite.dev group at GRC, no, no, no, no, no, do not put something that wipes data in a program that restores data. And that's like, oh. How about if I made this green, then? How about, like, red and flashing neon? No, no, no, no, no, no, no. I said, okay, fine. If you're going to make me charge more for it, then I guess I will. So...

Leo: How much are you going to charge, do you know?

Steve: Maybe 19 bucks?

Leo: Okay, yeah.

Steve: And it'll live forever and run. And you can run it on all your drives and so on and so forth.

Leo: Very nice. Very nice.

Steve: Yeah. And, I mean, again, you can use the free one from DBAN, or you can use Beyond Recall. And I don't even know what it's going to do yet except a perfect job. Which means understand about the relocated sectors and understand about drives that support their own low-level formatting or their own ATA wiping. And whatever it is, it'll do the Cadillac job of wiping a drive. So everyone has my guarantee of that.

Leo: How cool that you're doing that, thank you.

Steve: And there is a chart that came out. I tweeted this today. Backblaze, the big cloud storage folks, they just blogged today a blog posting: "What Hard Drive Should I Buy?" I imagine if you just google "what hard drive should I buy," that'll probably take you there. And scroll down to that bar chart, Leo, because I've got it in the show notes.

Leo: The annual failure rates.

Steve: Yes.

Leo: It is way lower on...

Steve: Hitachi.

Leo: ...Hitachi's. Seagate's way higher.

Steve: Yeah. And in fact they're seeing, in one case, 120% annual failure rate on Seagate drives, meaning they don't even last one year.

Leo: Wow.

Steve: Yeah, the ones...

Leo: So Western Digital's kind of in the middle and fairly consistent. But Hitachi's way down there, below 2%.

Steve: Yeah, now, the bad news is that Western Digital bought Hitachi. So we hope they leave them alone. But that's apparently - Hitachi at the moment seems to be the sweet buy. They're a little more expensive. And I mentioned, too, I saw a bunch of dialogue about warrantees, that warrantees are important. And I had also mentioned already that, yes, that's true, and that the drive manufacturers very quietly snuck the warrantees down. They used to be three years. They're now one year. And so they were realizing, wow, we're losing too much revenue because our drives are no longer lasting three years, which is a chilling fact. And it's a reason I'm selling SpinRite to this day.

Leo: They point out that they haven't had any Toshiba or Samsung, or enough Samsung and Toshiba drives, for good statistics. They do have quite a few. They have 12,000 Seagate and Hitachi drives, almost 3,000 Western Digital drives, and then a handful of Toshiba and Samsung drives. So this is probably statistically relevant. Although I would point out that most of these drives are more than a year

old, in some cases two and more years old. And so as a result what you're buying today may not be the same thing by any means.

Steve: True. Also it is the case, and they mention this, that they generally are doing write-intensive work, not read-intensive work. Whereas most users use of drives is the reverse. It's read-intensive and barely writing. So that's also worth keeping in mind. And the drives are never being powered down. That's actually the way I run all of mine. My servers are never powered down, and my own workstations here are never powered down. Drives really seem to last a long time if you just leave them alone and let them - don't get them too hot, and you let them just keep spinning.

Leo: We've got questions, nine of them. We're going to get to those in just a second. Yes?

Steve: Yep. I just wanted to give people a little update on SQRL. Code exists now. I will, by this time next week, it will be online and available for download if people want to play. It's not SQRL code yet. This is - because, again, I'm marching forward, laying down a foundation of technology as we go. This is the password-encrypting technology that deliberately takes five seconds on a smartphone, or maybe 30 seconds if you're wanting to export a key, like your galactic master identity key, where you want - obviously you want to keep it safe. And you're exporting it so that it's going to be safe, but you want to put it under a password which cannot be accelerated.

And so we do a couple things. First of all, we needed a process that could last 30 seconds. And it turns out there weren't any. I mean, nowhere is there a means of having something take 30 seconds. No existing technology does that. So we have developed it because that's what we want. We want something so resistant to brute force attack that a brute-force attacker, no matter how much technology and desire, and I mean even the NSA, cannot crack this.

So on one hand we want it to take a long time so that every single guess anyone makes, even you, even when you put the right password in, you're going to have to sit there and wait 30 seconds. Well, that's not a problem if it's for importing your offline-saved galactic master password. But even when you're authenticating to your phone, the idea is, remember, that SQRL provides slam-dunk replacement for username and password. But you still need to prove it's you using your phone. So there our feeling is, eh, five seconds, if you only have to do it, like, whenever you run SQRL, that's not too long to wait for the security of a bad guy having to wait that long to guess those passwords.

So what we did was we created this notion of using Scrypt, S-c-r-y-p-t. That was the technology that Colin Percival developed for his Tarsnap multiplatform cloud backup solution. We've talked about that. Tarsnap is a very good multiplatform solution for allowing people to do offsite storage, powerfully encrypted. He wanted something better than Bcrypt, which is all that was there at the time. And of course we've talked about PBKDF2, Password-Based Key Derivation Function 2. There's an RFC for it. It has an ITF standard. It iterates, but unfortunately it uses typically the SHA-256 hash. Well, that's the hash that all of the crypto currencies are using, so people have hugely accelerated the speed of that hash. You can easily find FPGA code and even now, of course, ASICs, which have been custom designed to do SHA-256 at light speed. So basing an iterative password-based solution on SHA-256, unfortunately, allows it to be hardware-accelerated much too quickly.

So what Colin did was he created this notion of a memory hard function that we talked about, where you use a pseudorandom function to fill a large area of memory. And our standard is 16MB because no GPU or FPGA or ASIC has fast access to 16MB. They may have local caches - and same thing for a regular CPU. They may have local caches that give them access to a fraction of that. But the way this thing works is it uses the random data in the 16MB to choose the next access of somewhere in the 16MB. And the data there chooses the next access in the 16MB. So it all has to be there at once. And there's essentially no way to spoof that. So what this does is this thwarts any attempt at using hardware acceleration. But even Scrypt, when you're using those sorts of parameters for 16MB, it doesn't take long enough.

But it turns out you can chain Scrypt. So what we do is the user's password is put in, along with a random salt, into the first instance of Scrypt, which probably runs, in the benchmarks that we've been doing, it runs maybe a 30th of a second. It results in 256 bits, or 32 bytes. We use that as the seed for the next one, and again the user's password, and run that. And then use the output of that as the seed for the third instance, and so forth. So we iterate over Scrypt each one of these instances requiring 16MB, and the output of that one being the input to the next one. Then we end up XORing all of the outputs which occurred, and that gives us our final 256-bit key, which is derived from the user's password. And as far as we know, there is no way to speed it up.

And the other thing that's so cool about this solution is that you can run it either for "N" number of iterations or for "S" number of seconds. So you can say, give me a 30-second encryption of my password, and it iterates watching the time pass until that length of time has occurred and then says, here you go, and here's how many iterations that was. Because, when you're decrypting it, you need to use the same iteration count as you used when you encrypted it. But so we end up with a system for SQRL which allows you to either specify time or iterations. It's memory hard and no way to speed it up. And this is brand new crypto technology. Nothing like this exists in the industry now.

The project I am on when I finish this podcast today is getting this documented. And I do have a very cool Windows console app which has been well vetted. Two other people have reimplemented the algorithm that I explained verbally in the newsgroup, in different languages, and they generated exactly the same output given the same input. So we have verified three ways that we're all in agreement about making this thing cross-platform. And one of them, by the way, is someone who's doing SQRL for iOS. And so it'll be available there. And somebody else is doing it for Android. So next week I'll have some URLs for people.

But anyone who has Windows or Wine on any of the UNIX systems can run this, and it serves as an interesting benchmark because this shows how many iterations your computer requires. Or you tell it you want a hundred iterations, and it shows how long it takes to do a hundred iterations. So you're able to compare, essentially, the memory bandwidth and throughput of your machine. It's very cool.

Leo: Can't wait. Can't wait. Let's get to our questions. Are you ready, Steve Gibson?

Steve: I am. I did want to just mention a bit of errata. Many people commented that I was calling COTS, the acronym for Commercial Off the Shelf, I was saying "common off the shelf." So I certainly stand corrected. I know the difference. I just got started on the wrong track, and I kept saying the wrong thing all through the podcast.

Leo: Happens all the time. To me, not to you. Here we go, Steve. We've got some cooking in here, as well. Just watched the coffee-making - this is from Guillaume Auclair in Sherbrooke, Quebec, Canada. Just watched the coffee-making episode that you and Leo did on New Year's Eve. Did I not get the ratio of beans per hot water, ounces of bean versus ounces of water? So Steve has his recipe here. This is the key.

Steve: Yeah. We're going to do this again next year. But so that people don't have to wait, I did some measuring so that I could give people what the ratio is. So I take what is a quarter-cup dry measure, 60ml is also written on my little quarter-cup scoop, of whole Starbucks Espresso bean. And that's a level quarter cup, not heaping. When I've made it heaping by mistake, the coffee's been a little too strong for me. So just sort of a flat level quarter cup of beans. So then I grind them for drip brewing, which is a rather coarse grind, using a good burr grinder. Then they are drip brewed through a brown paper Melitta-style paper filter. And the input is about 750ml of clean reverse-osmosis-filtered drinking water. So that's the ratio, 750ml of water through a quarter cup of drip-ground beans, filtered through a Melitta-style filter. And, boy, I mean, I've had people stunned by how good this coffee is - smooth, rich, and creamy. And people who normally need cream and sugar for their coffee just don't need it with this. So that's the recipe.

Leo: Here's a tweet from @daveheld_info. Seems a bit confused, he says, about the NSA ANT coverage: If the NSA can do this in 2008 - and those were the ANT slides we saw, five years old - who else can do it now? Aren't you worried? Is it all a big joke, like the podcast was? I don't know what that means.

Steve: Well, okay. So I think Dave is upset with me that I'm not more upset about this. And that's sort of my nature. I mean, I don't get upset about things I don't have any control over. And I have no control over this. So I don't think this is a joke at all. I think this is worrisome. But on the other hand, none of this was mass surveillance technology. This was old school, bug somebody's office because you need information from them, and irradiate their bug with a cool radar beam in order to get that information. And, yes, this is old technology. But this is what I would hope our taxpayer dollars are going for. I mean, we want our intelligence gathering to be able to do this kind of job. So I'm not upset by this because I'm here to report it and to sort of explain the technology. And again, to me, this seems an appropriate use of these kinds of spying resources.

Leo: Yeah. I guess you could have the attitude that nobody should be allowed to spy on anybody else. But good luck.

Steve: Well, yeah. Good luck with that.

Leo: No spying allowed. I don't know. I don't think that's an unreasonable point of view. But it's not practical in this day and age.

Steve: And it's not going to change the world.

Leo: No. Bill Gearhiser, Boca Raton, raises a common question about VPNs. He says he signed up for proXPN, heard about it on the show: I wonder if Steve's evaluated the issue of identifying the far end of the pipe on proXPN. For example, if I connect to Amsterdam and start browsing, I know that's fairly anonymous. Now, we've got to explain, a VPN is not about anonymity, but all right. But if during that session my mailer does a poll for email, won't that poll pop out at the same end of the pipe to Amsterdam? Won't that identify who owns the pipe? If so, what does one have to do to anonymize a bit better? VPN has never been about anonymity.

Steve: Yeah, and that's why I chose this question, because it is a common misunderstanding. The way to think about a VPN is that it protects your local environment. That is, without that at Starbucks, you are really not safe because Starbucks is open WiFi that is unencrypted, and anybody sitting there sniffing the traffic gets anything that's unencrypted. So, which typically is like all of your email, unless you're explicitly encrypting, for example, using Google Mail now finally is over HTTPS connections. But typical POP email is still often an IMAP or still often not. So the idea is you use a VPN to get out of your local environment, whether it's Starbucks or the hotel that you're in, which is another common source of attack, or even your ISP that increasingly is, unfortunately, not very trustworthy with your own traffic. You want to get it away from them.

So it is absolutely true that, when it emerges on the Internet at one of these aggregation points, where the VPN server is that you've connected to, then essentially it's out of the vicinity that was in danger, and now you're just sort of out in the middle of the Internet somewhere, probably without anybody looking at it, although you could argue that VPN servers are other targets of opportunity, sort of in a variant of the way the Tor servers, the Tor exit nodes represent a target of opportunity. But...

Leo: I guess the difference is it's the difference between local anonymity and global anonymity. You are locally anonymous, but you're not globally anonymous. And so 8bitsteve made the point in the chatroom, it does say in the proXPN copy, surf the web anonymously. And I haven't read the copy, but we should probably explain that that means, from the point of view of your ISP or anybody sitting next to you in a coffee shop, not globally anonymous. And as it turns out, global anonymity is challenging even with Tor.

Steve: Yes, yes.

Leo: It's a tough thing to achieve.

Steve: It's really not something that the Internet offers. It wasn't designed...

Leo: Yeah, that's kind of my point earlier on is that we're asking an awful lot of the Internet when we say we want to do everything privately.

Steve: Right.

Leo: Eric in Santa Cruz, my old stomping grounds, wonders about an IPv6 version of ShieldsUP!: Curious when you plan to set up ShieldsUP! over IPv6. I think it would be a great addition. Do you have any plans for that?

Steve: So I would love the time to do that. I actually did purchase some hardware that would begin to take me in that direction. I talked to Level 3, my main pipe provider at GRC, and also to Cogent, that runs my T1s, because I would need my T1s to have IPv6 space in order to develop the technology here, which I would then carry to Level 3 and install at GRC. All I need is time. I would love to rewrite for v6. But as everybody knows, I have a bunch of things on my plate at the moment. I've got to get SQRL launched. Then it's back to SpinRite to get 6.1 launched. Then it's, I think, Beyond Recall, to get that done.

Then I'm really thinking, the way the world has gone, I should go back and take a look at CryptoLink, but do it as open source software, never intending for it to be commercial. But we'll sort of play that by ear. We'll see what the world looks like when I'm on the other side of SpinRite 6.1, and everybody's got that, and I've probably got Beyond Recall finished because I think that needs to follow. Then we'll sort of see where we stand. All I need is time, and I'm the only one writing the code, and I'm sort of - I'm a turtle. I'm slow and methodical, but the stuff lasts a long time. So it's just a matter of getting to it.

Leo: Jeff Levy in Poughkeepsie, New York, asks about TrueCrypt: I recently bought a 64GB flash drive, decided to do whole-drive encryption with TrueCrypt. While going through the process, I noticed all of TrueCrypt's algorithms are 256-bit. Does that mean my drive has only 256-bit encryption, even though my password is 30 characters long? Or does the strength of the encryption lie in the password length? Love the show. Long-time SpinRite user.

Steve: Now, you know, I'm not sure, if you chain encryption algorithms, they almost certainly use 256 bits for the various key lengths. I don't know if the other encryption algorithms - because I didn't do any research on this specific. This just occurred to me as I was listening to Leo read the question. If you use AES, which is the standard, then you could use it with 256 bits. And it is, 256 bits, I mean, that's now, as I have said in the past, 256 bits is the new black. It is all anybody needs. It is your absolutely private bitcoin identity that no one can guess. It is your BitTorrent Sync ID that no one can guess. It is your master key for SQRL, my project of the moment, which no one can guess. I mean, it is impossible to guess that.

So what they do is they use your password to encrypt a randomly chosen 256-bit key. When you're setting up TrueCrypt, and they ask you to, like, move the mouse around in random directions in order to, like, add additional entropy to what they've already got, that's the pseudorandom, 256-bit master key which is used to encrypt your drive. But then that is encrypted under your password. So the way someone would crack your TrueCrypt-encoded volume, they would never, never try to start guessing 256-bit keys. There are just too many of them. They would always guess passwords, run it through TrueCrypt's algorithm to decrypt the password into a candidate 256-bit key, and then see if that works. That's the way you crack it.

Now, it is likely, because I'm sure the TrueCrypt guys did this right, that if you chain cryptos - remember that there's an option in TrueCrypt of not only using just AES, but also using Blowfish or maybe 3DES and who knows what, if you chain - and you could,

like, use them all. They probably take another chunk of entropy for the key for the second one and another chunk of entropy as the key for the third and so forth. So you actually do get a longer key. But to my mind, that's just crazy overkill. The only reason you would want that is if a defect were found in the crypto you were using, having them chained would protect you under the other ones. But AES is really living up to its reputation, and no one is finding any problems with it, despite looking really hard.

Leo: Good.

Steve: Yes.

Leo: Dan Murfin in London, U.K., notes another remote access router disaster: As a customer of EE in the U.K., and someone who uses their Bright Box 2, I was a little dismayed to learn this morning that it is subject to yet another remote access vulnerability. My question is can we really trust any of these ISP-provided routers? Or should any security-conscious user buy a third-party router they can fully configure?

Steve: I think, given what we have just seen with this port 32764 disaster - and you know what, Leo, I'm feeling, like, guilty that I didn't suggest I jump onto your...

Leo: Well, we talked about it. We talked about it on the show.

Steve: ...weekend show. Actually, I thought you had because I saw some traffic patterns on the weekend where I thought, oh. I looked at the clock when I saw the site...

Leo: Realized who was on.

Steve: I thought, okay, Leo must have just talked about this port because - yeah.

Leo: Yeah. I was talking about securing a wireless router. I do this talk fairly frequently, you know, the five things you need to do. Except now we have to add another one, which is test your - go to bit.ly/port32764 and test that router.

Steve: Yep, exactly. And so my feeling is, given what we've just learned, anybody who is concerned about their security ought to get a router that can accept either the Tomato or the DDWRT known clean firmware and load that up. Both of those projects produce a beautiful router that is feature complete and - just reflash a router that can accept it, and then you know what you've got. Otherwise, look what we keep learning. Not long ago it was the open UPnP port disaster. And now we've got this obscure random port up in a high port space. Wow. I just think, just switch to some firmware that you can trust because this is your Internet-facing self. This is what the world has of you is your router.

Leo: Do you like OpenWRT? What is it that you like the best, Steve?

Steve: I'm in big-iron routing, so I'm not using any. But I always say either OpenWRT or the Tomato firmware, both.

Leo: And those are open source, so people keep them up to date and avoid these kinds of...

Steve: Yeah. And there's a great team in both cases that are maintaining those.

Leo: Tom Walker, Littleton, Colorado.

Steve: Speaking of the devil.

Leo: Speaking of the devil. He's wondering about port 32764, as well. Steve, a quick question. Wouldn't either my Windows Firewall or my antivirus firewall block port 32764? For some reason, he says, this reminds me of the first program I wrote back in 1979: Let X=0, X=X+1, PRINT X, GOTO 20. In other words, have the computer count up by one. He said: I wanted to see how long it took a computer, a TRS-80 Model 1 with 4K of memory, to count to a million. It never went past 32767, then errored out with "Overflow occurred at line 20." So is it a coincidence? Is 32764...

Steve: Actually, they're closely related. The fact that his TRS-80 crashed out at 32767 tells us that that version of BASIC was using integer math with 16-bit variables because - oh, and it's signed, by the way, because 32767 is one less than half of 65536. And the way two's complement math, which is what this is using, the way it works is the high bit, the most significant bit is considered the sign bit. So if you force that to be zero, meaning the number is positive, and all the other one bits are on, then that value is 32767. So the maximum positive value that a 16-bit signed value can contain is 32767. So when his BASIC program was at 32767 and tried to increment that, it overflowed its representation. It could not represent 32768, the next one up, in a 16-bit signed register. And so it just exploded.

And as I mentioned before, 32764 is four down from the midpoint of the port range, which goes from 1 to 36765. The midpoint is 32768, and 32764 is four less than that. So, yes, they are related. And the answer to Tom's question, which we sort of answered just in the prior one, the problem with that port 32764 is not Windows Firewall or Windows or anything inside your network. It's the router itself. It's the router that's on the front line. That's where there is a service, an obscure little backdoor trojan, really, which is listening to that port and accepting commands. So nothing inside your network can help you because it's the router itself that is going to accept that connection and execute commands from a bad guy.

Leo: Let's see. Andrew Stenenson, Dorset, U.K., muses about XP updates in our next question: Should Microsoft stop all security patches for XP? Or maybe should

they have done it more gradually? I don't know how you do it gradually. What I mean is it seems rather brutal - brutal - to pull the plug on all security patches all at once. Maybe they should continue to patch the most severe remote code execution vulnerabilities until XP usage has dropped to an arbitrary percentage. Slightly good news for XP users: Microsoft has extended update support, he says, for Security Essentials until July 2015. That's not quite true. They're going to update the virus definitions but will not update Security Essentials. Big difference.

Steve: Yeah. So as an XP user, I guess the only thing - there's a number of things happening here all at once. First of all, the thing that's annoying is that Microsoft is essentially getting XP security fixes for free for the most part because all the same things, what we keep seeing is when there are vulnerabilities found, they exist in Windows 8 and 7 and Vista and XP because it's a common code base. My complaint with this constant version churn is that it's not for the users' sake, it's for Microsoft's sake. It's for upgrade revenue, largely. I mean, look at what a catastrophe Windows 8 has been. Why did they force everyone off 7? Well, because they can get revenue from upgrading people.

So I guess what's annoying is, if there was a security vulnerability that only affected XP, I could understand them not patching it. But they're actually addressing vulnerabilities in the core shared code of all of these OSes. So why not toss in XP if it applies, rather than just saying no, from April 8 on of 2014, you don't get patched. Because what that does, of course, is it's forcing upgrades of people who don't want to upgrade. Their XP is working just fine, like mine is. I have no interest in upgrading. But I'm not - I and a world of other people are, what is it, it's like 43% of Windows is still running XP, something crazy like that, because it's just fine. But Microsoft is saying, okay, we're not going to continue patching, even when we have the patches. Even when we develop them, we're not going to give them to you because we're going to make you upgrade. It's like, eh, okay, well...

Leo: It has been 13 years, Steve.

Steve: But Leo, Eudora's working just fine for me.

Leo: They haven't updated that, either.

Steve: I know. They're dead and gone, but it works just fine. There's this weird, weird mindset that people have developed that new stuff is better. Look at Windows 8. I don't think that's any better. Paul, didn't Paul declare it a complete disaster?

Leo: I don't think Paul said that, although I keep trying to get him to. It's pretty - it's a disaster, all right.

Steve: Yeah, talk about keeping yourself busy just explaining to people how to use it.

Leo: This next one seems to come from the same person, Andrew Stevenson in Dorset, U.K.

Steve: Oh, it's weird that I - it's funny, too, because when I was putting it in, I thought, wow, we've got a lot of people in Dorset, U.K.

Leo: Yeah, same guy, I think. So...

Steve: So sorry to everybody else. I didn't mean to be giving Andrew excess time.

Leo: I think it's a typo. But maybe not. So you're saying this actually is the same guy?

Steve: It probably is.

Leo: Steve and Leo, there's been a change to the way that SSL Labs rates TLS connections. Congratulations, Steve, you're now an A+ with extra credit because you use Strict Transport Security. If you want to attain the green bar on SSL Labs for using forward secrecy ciphers, I believe you need to add the older DHE cipher suites to your server, which would change your rating from "with modern browsers" to "robust," which you can see down in the protocol details. I guess Ivan Ristic, who writes these tests, has written about what he's changed. This is an ever-changing standard, in a way; right?

Steve: Well, but this is significant because SSL Labs is a site that I know our users are huge fans of - it's ssllabs.com - because you can put any server into it, like Google or Microsoft or GRC, dotcom in every case, and it'll tell you about the security suites that the server offers. To get that security, your browser, as we know with SSL, has to be able to have compatible security suites. But this will show what the server offers. The big change here is Strict Transport Security is not SSL, it's HTTP.

So what has happened is SSL Labs has dropped to the, depending upon how you look at it, the layer above or the layer below, probably the layer above, at the protocol application layer, and it's noticed that GRC, when you make a query from GRC over SSL, or not, actually, it doesn't matter, although GRC now forces SSL, or if you try to connect without it, we redirect you to the same page over SSL, so that we've upgraded all of our links and so forth. The point is that, after your connection, anything, any asset you request from GRC, we send back the Strict Transport Security header with a long expiration. It's, I don't know, what is it, 301 something something something. It's like basically I'm saying, forever use only SSL. So what that does is browsers will cache that, and browsers will then silently upgrade the HTTPS connections themselves to - I'm sorry, the HTTP connections to HTTPS, because they have cached permission to do that that they received from GRC for that long period of time. And so we're now rating an A+ at SSL Labs because we're doing that.

And of course, as our listeners will remember, I went one step further and asked Adam over at Google to build that knowledge into Chrome so that even the first time, there's

one little tiny remaining exploit possibility, which is if a browser had never visited GRC or didn't understand about Strict Transport Security, the first connection could be over HTTP. Not with Chrome. Chrome knows, it's built into Chrome that GRC.com will always be HTTPS. And so even the very first time you connect, Chrome elevates the connection, upgrades it to SSL, even if you didn't ask for it. So this is a nice step forward for SSL Labs. They're making their tests more comprehensive by going past the SSL handshake into the actual HTTP protocol and looking at what the server does. So that's very cool. So thanks for bringing that to my attention, and now our listeners'.

Leo: Yes. Well, my friend, we have run out of time. No, we haven't, we've run out of questions.

Steve: We've run out of run out.

Leo: Run out. We're all out. But it's always fun to do this. If you have a question for Steve for future episodes, don't email him, just go to GRC.com/feedback. That's his feedback form. And I guess it does, I didn't know this, but it does go to his email box, where it will join 50,000 other questions.

Steve: It goes actually to a separate, like off on the side, Security Now! email box. And then I normally pull it when we're doing a Q&A to get the latest. And I sort through those, browse through them, and find good things.

Leo: Very nice.

Steve: And the good news is it's not piling up forever any longer. Now I've got MailStore to send them all off to.

Leo: When does, what is it, Total Recall, Total Non-Recall, go on sale?

Steve: Oh, Beyond Recall. It'll be after SpinRite 6.1.

Leo: Okay. So SpinRite is there at GRC.com, SpinRite 6.0. Free upgrades for life to 6.1.

Steve: Free upgrades to 6.1.

Leo: You'll also find a lot of freebies. Steve's very generous with his time, and a lot of security updates and things like that. And 16Kb versions of the show, which Steve edits with his very own hands using a razor blade and a grease pencil. He also makes transcripts, handwritten transcripts available, from Elaine Farris, at the same place, GRC.com. We have somewhat larger, automatically edited versions of audio and video available at our site, TWiT.tv/sn. And of course you can always subscribe

after the fact, anytime, at all of your favorite podcatch clients. We do the show, I said Wednesdays, Tuesdays now, 11:00 a.m. Pacific - no, I'm sorry, 1:00 p.m. Pacific, 4:00 p.m. Eastern time, 21:00 UTC. That old time creeps in.

[Talking simultaneously]

Steve: Old habits die hard.

Leo: ...long time. Long time. And you can watch live. We love it if you do. Otherwise we'll see you on the Internet. Thanks, Steve. We'll see you next time on Security Now!.

Steve: Thanks, Leo.