Listener Feedback 217

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

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-522.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-522-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to talk about the week's security news. And then, finally, it's been slow enough in the aftermath of Black Hat and Defcon to actually do some Q&A. So we've got 10 questions, 10 answers, coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 522, recorded Tuesday, August 25th, 2015: Your questions, Steve's answers, #217.

It's time for Security Now!, the show where we protect you and your privacy online with Mr. Steve Gibson, the king of the tinfoil hats. But that's a good thing. That's a good thing. You want somebody who's really paranoid to be running a show like this. Hi, Steve.

Steve Gibson: Well, hi, Leo. It's great to be with you again. And I, you know, I'm the first to say that what interests me is the theory of what can be done. And I know from the great feedback we have from our listeners that that's our audience.

Leo: They appreciate that, yes.

Steve: You know, people who absolutely want to know. And it may be that we choose not to worry about it. And, for example, you and I were just talking about messaging. I use iMessage, knowing that Apple is the curator of the keys. Which means, if I were doing anything other than arranging when I'm going to have a meal with a friend, I would be using something truly secure. But it's not a need that I have. But I think, from a theoretical standpoint, it's interesting to talk about the technology, which is of course what we do on the podcast.

Leo: Also, it's good to know, if you did need to communicate securely, that this is not the way to do it.

Steve: Correct.

Leo: I think that's worthwhile; right?

Steve: Absolutely.

Leo: Yeah.

Steve: Absolutely. Especially in this era where it's no longer a question whether there are forces outside of ours that would love to know what we're talking about, just because they'd like to. I mean, they're like, well, you know, we don't, I mean, it's in the wind. We don't think anybody should be able to have a secure encrypted conversation because bad guys shouldn't; and therefore, since we can't tell the bad guys from the good guys, we need to be able to listen to everybody. Which is, I mean, that is the prevailing philosophy in law enforcement today. So if somebody in a free country objects to that philosophy, then the degree to which they want to protect themselves, we can give them everything, you know, a spectrum of choices, from communicating in the clear to truly bulletproof privacy. Yeah. And so I think that's great.

Leo: So believe it or not.

Steve: Yes. Finally.

Leo: Believe it or not.

Steve: In fact, I led the show with a tweet from our friend Matthew Green at 7:41 a.m. on August 25th, today. Matt, of course, is the cryptographer we're often citing who did the second round of - actually managed both rounds of the TrueCrypt audit. He's the Johns Hopkins cryptographer at Black Hat and Defcon and so forth, very much involved in what's happening. And he said: "It's almost as though people blew all the interesting security results on Black Hat and Defcon, then took the rest of August off."

Leo: That's exactly what happened.

Steve: Yes, it is exactly what happened.

Leo: Of course.

Steve: Yeah. There was all the pre-show leakage of stuff that's like, oh, my goodness, you know, the Android problems with StageFright, and then the Jeep Chrysler problems and so forth, leaking out before. Then a bunch of really interesting topics during the show. Now, quiet.

Leo: Yeah.

Steve: And so the point is we finally get to do a Q&A.

Leo: Woohoo.

Steve: I wanted to talk a little bit about Lenovo's BIOS behavior.

Leo: Oh, good.

Steve: They've had a retraction and update since we discussed them last week. Believe it or not, there is an open source ransomeware file encryptor on GitHub. It's like, okay, what world are we living in? Many changes coming to Firefox. I want to discuss a little bit about the consequences of the growing intersection of life and the Internet; what I regard as the best thing about Google's new router; something we've never really talked about much, the need for physical security; and some miscellaneous stuff.

And then 10 questions and thoughts and observations. Of course lots about what we've been talking about recently because everyone really is interested in the virtual machines and Sandboxie and adblocking and so forth. So a great mailbag from our listeners, beginning to catch up. I had 630-some pieces of mail when I hit the refresh button this morning.

Leo: And not one word about Ashley Madison.

Steve: Actually, that's the consequences of the growing intersection of life and the Internet.

Leo: Amazing story. But there's nothing to say except, oh, well.

Steve: I have an angle.

Leo: All right, Steve. Time for the news segment.

Steve: So we covered the Lenovo problems last week. And also in general this whole issue of sort of rootkit-y behavior on the part of BIOSes, the idea that a BIOS could write files into the file system. Well, that is not the formal accepted practice by Microsoft. They've got that autochk.exe file in their Windows\system32 directory that we talked

about last week. They don't want somebody replacing that. You can imagine that didn't go over very well in Redmond. So Microsoft does have a technology which they do support, which I described last week, which uses the ACPI functions in the BIOS. There's a way that the BIOS can configure some ACPI tables to tell Windows where in main memory the BIOS has just copied an executable file. And then when Windows starts up, it reads the table, goes up into physical memory, and then copies that itself into its file system. And Microsoft is happy with that whole approach.

Well, this, it turns out that the way Lenovo had implemented theirs was, as we know, aside from being sort of making people feel very uncomfortable and seeming sort of underhanded, they got caught doing this, using this to install stuff that users specifically didn't want installed. So Wednesday of last week, the day after we described all of this, they put out an update. And they have what they call the LSE, is the Lenovo Service Engine.

And so they said: "Lenovo's use of LSE was not consistent with Microsoft's guidelines. As a result, LSE is no longer being installed on Lenovo systems. It is strongly recommended that customers update their systems with the new BIOS firmware which disables and/or removes this feature." So this is not, apparently, being made available in an auto update mode. They have removed it from their products going forward. What I did was many people after last week's podcast wanted to know what specific model numbers were affected. Of course, that's natural. It's wait a minute, is it my laptop that is doing this?

So I made a bit.ly link for this 522nd Security Now! podcast, bit.ly/sn-522, all lowercase, bit.ly/sn-522. That will take people to their announcement. And then, if you scroll down a bit there, Leo, there is a list way too long for me to read into the podcast. And since it doesn't affect the majority of our listeners, I would imagine, I didn't want to bother everybody with it. But it is a lengthy list of specific model numbers of laptops. And embedded in the text toward the top of that are rather obscure links, but they're there, to the desktop firmware update and the notebook firmware update.

Anyway, I do have all of these links in today's show notes. So anyone can just grab the Security Now! 522 show notes, and right up at the top you will find links to those updates.

Leo: Why is it, you think, that they didn't do all of their computers? They only did, for instance, I don't think the ThinkPads are on this list. The X laptops aren't on this list.

Steve: Yeah. It seems to be the higher end, sort of the ThinkPad-ish generation laptops. And I think it's, frankly, it's because of crapware. I mean, you notice that they're saying...

Leo: Superfish wasn't on the high ends either. I mean...

Steve: Right, right. So they call it the Lenovo Service Engine, but we've never really had a clear explanation of exactly what service it provides. And the sense is it's not a service for us, it's more of a service for them, you know, because the presumption is they were putting stuff on that people didn't really want, and that this was about keeping it there.

Leo: And maybe a higher end audience would have noticed it more; right?

Steve: Very likely.

Leo: Maybe the IT department who buys a lot of ThinkPads might have figured this out.

Steve: Oh, well, especially if you reinstall Windows, and then this thing modifies your installation. You know, that's what crosses the line, the idea that somebody could wipe their hard drive or replace their hard drive or do a fresh install, and this thing reaches up out of the BIOS before it starts up Windows...

Leo: Horrible.

Steve: ...and changes files. It's like, eww.

Leo: Horrible.

Steve: Yeah.

Leo: But, yeah. Now, what do we think? Are we not going to recommend Lenovo anymore? I mean...

Steve: Eh, I mean, okay. So it's - clearly the damage has been done. I tweeted that link this morning, just so that the people who follow me on Twitter had it. And enough people took the time to send back, you know, they're dead to me now. Trust has been broken. First Superfish, then this. They clearly don't have their interests, best interests in mind. You know, I've got two Lenovos in good shape, an X61s, I think it is, and something else, that I love. And I'm hoping that I'll be able to get...

Leo: Great computers.

Steve: That I'll get replacement batteries for them as long as I need them. So, yeah.

Leo: Well, notice IBM is no longer using Lenovo laptops. They've switched to Macintosh. I don't know if this is related or what; but, yeah. It's sad. It's really sad.

Steve: Yeah, well. And...

Leo: They own Motorola phones, too, which worries me a little bit because I love the Motorola phones. But...

Steve: I think the reality of market pressures is that there is just - there is so much pressure on the laptop makers to generate revenue through additional channels besides the consumer. And so what they do is basically they're subsidized. It's very much like, you know, smartphones. You don't buy a smartphone for list price. You get it subsidized under a contract. And similarly, the only model that works in a laptop is if they put stuff on there that they are paid to put on by third parties.

And it's always sold as a benefit, you know, I mean, like McAfee or Norton or Symantec tends to be on these things. Those companies are paying for that stuff to be on there. And of course it'll go for a year, and then you start getting expiration notices, and they want to convert you over to a paying customer. So, unfortunately, the economic model has become - we can't charge what we need to, or then we don't look competitive. So we're basically selling subsidized laptops by adding stuff that we're paid to add.

Leo: This has been a problem for so long, and even Microsoft tries to fight it.

Steve: Yeah.

Leo: But it's some, you know, I hate to say it, but a little bit our faults for pushing for low-priced laptops.

Steve: Right.

Leo: When you cut the margins to nothing, they've got to make it up somewhere.

Steve: Right, right, yeah. So I think it was our friend of the podcast, Simon Zerafa, who pointed me to a GitHub link, and I just - okay. So it's "hidden-tear." So I think you can probably just google maybe "GitHub hidden-tear," with a hyphen between them. So this is - it's a ransomware file encryption sample which can be modified for specific purposes, says the description. And it has a bullet list of features: uses AES algorithm to encrypt files. Okay, good, state-of-the-art encryption. Sends encryption key to a server. Okay, that's what we want from a - that's the way you do ransomware properly. Encrypted files can be decrypted in decrypter program with encryption key. Okay, yeah. Creates a desktop file in desktop with a given message. So that's where the malware leaves its explanation of what to do. Small file size, 12K. Doesn't detected - it actually says "Doesn't detected to antivirus programs," so we know this is not an English-speaking person who put this up.

And that was as of August 15th, so, what, 10 days ago it was unknown by any AV. And it says in the benefits: "Target file extensions can be changed." And then it provides a default list, which is a comma-separated text string of file extensions - .txt, .doc, .docx, .xls, .xlsx, .ppt, .pptx, .odt, .jpb, .png, .csv, .sql, .mdb, .sln, .php, .asp, and so forth. So it says: "Legal warning. While this may be helpful for some, there are significant risks. Hidden-Tear may be used only for educational purposes. Do not use it as a

ransomeware! You could go to jail on obstruction of justice charges just for running Hidden-Tear, even though you are innocent."

Leo: They even have a YouTube video, demonstrating how it works.

Steve: It's a full-feature, full-service, open-source ransomware.

Leo: It's hard to believe.

Steve: I know.

Leo: Apparently there are several other Hidden Tear projects. I don't know if they're forks or improvements or what.

Steve: The good news is that somebody who cannot create this themselves, I mean, there's a reason this is here, which is, unfortunately, the prevalence and pervasiveness and zero cost of state-of-the-art crypto makes this simple. This is, I mean, I could write this in an afternoon. I mean, anyone, any of the coder crypto people could. This is a trivial problem to solve. And what's so damning about it is that it's so effective.

Now, the good news, though, is to be actually useful, you need two things. You need an infrastructure, I mean, the whole payment system. You need a server to receive the encryption keys. You need to obfuscate that server. Actually, you need a network of servers with rolling pseudorandom-generated domain names so that your cryptoware knows where the server will be, on what domain name, at what date and time in the future so that, as these domains get taken down, they're being constantly replaced.

I mean, and you need a way to get this into a person's machine. I mean, if you were purely malicious, you could arrange to run this on, you know, send this as an EXE to, well, not a friend, someone you really dislike. And if they're dumb enough to run an executable that they receive, then that would be bad. But there's a long way to go from an open source 12K thing that basically randomly generates a public and private key pair, sends the private key off to a URL, then uses the remaining public key - oh, and then wipes the memory, hopefully, if it's any good, I'm sure this thing doesn't bother - then uses the private key to do a front-to-back search for files by extension and generates a per-file random key, which it then appends to the header and encrypts the file.

I mean, again, we've talked about how this stuff works. Trivial to do. But the good news is that's the least of turning this into a money-making enterprise. And anybody who's capable of actually getting into other people's machines on a massive enough scale to be useful, and creating a sophisticated money delivery infrastructure, and you've got to be able to have a communications channel with the users whose systems you've encrypted so that you are actually able to get them back the decryption key. Otherwise pretty soon people will realize no one's getting decrypted after paying you, and so it'll die because of your bad reputation.

The point is, you know, it takes much more to actually make money from this kind of ransomware than just dropping a sample out on GitHub. But still, I just had to sort of

shake my head. It's like, okay, wow.

Leo: Geez.

Steve: What's the world coming to?

Leo: Well, everything's there.

Steve: Yeah.

Leo: Information wants to be free unless it's encrypted, in which case you can give me a kebab.

Steve: Well, and exactly. And that's, you know, the information wants to be free. That's the biggest argument I have against the whole NSA, FBI, we need to be able to encrypt all communications because the world already knows how to do that.

Leo: Yeah, yeah. A little too well.

Steve: Yeah. And so all that they would do is force the law-abiding consumer...

Leo: Oh. Are you back?

Steve: Yes. The lights are blinking.

Leo: I saw them blink.

Steve: I stopped talking, and then the lights kept blinking, sorry. Like, wait, I'm...

Leo: I can tell by the blinking lights. Okay. We're not going to call you back. You froze solid, though, for, like, five seconds. That was weird.

Steve: Interesting. Yeah, the very last question of today is, so what's up with the audio, because I just want to make a mention to our listeners.

Leo: Just little hits in there from time to time.

Steve: Yeah. So, okay. So anyway, so the point I was making with the NSA is that the cat's out of the bag. Bad guys will use bulletproof, unbreakable crypto, which at the

moment everybody's using. But if it turns out that that's illegal, then they'll keep using it, and everybody else will just have law enforcement-breakable crypto. So it just doesn't make any sense. I'm hoping this thing finally gets resolved.

Leo, there's a very funny video. You've probably seen it because you're hip to all this stuff. It's got 12 million views on YouTube.

Leo: Wow.

Steve: It's a comedy. It shows a group of people, I think it's titled "The Expert," and it's a meeting where a client is asking a contractor to do something that's impossible. And the contractor has a couple, has brought along some domain experts, you know, specific knowledge experts. And they want nine red perpendicular lines, and some of them transparent. There it is. Yup.

Leo: Yeah, yeah, yeah, yeah. You want to listen to a little bit of it?

Steve: Sure.

[Clip]

FEMALE VOICE: In pursuit of these objectives, we start a new project which will require seven red lines. I understand your company can help us in this matter.

MALE VOICE: Of course. Walter here will be the project manager.

Leo: I've been in this meeting, by the way.

MALE VOICE: Walter, we can do this, can't we.

Steve: I know.

Leo: Who hasn't; right?

Steve: We all have.

MALE VOICE: ...our expert in all matters related to the drawing of red lines. We've brought him along today to share his professional opinion.

FEMALE VOICE: Nice to meet you. Well, you all know me.

Leo: We know who does the work here; right?

Steve: Uh-huh.

Leo: This reminds me a little bit of the...

FEMALE VOICE: We need you to draw seven red lines, all of them strictly perpendicular. Some are green, and some are transparent.

Leo: "Mr. Robot."

Steve: Yeah.

FEMALE VOICE: Can you do that?

MALE VOICE: No. I'm afraid that...

MALE VOICE: Let's not rush into hasty answers. [Crosstalk] carried out. At the end of the day, you are an expert.

MALE VOICE: The term "red line" implies the color of the line to be red. To draw a red line with green ink is, well, if it's not exactly impossible, it's pretty close to being impossible.

MALE VOICE: What does that even mean, impossible?

MALE VOICE: I mean, it is possible there are some people, say, suffering from color blindness, for whom the color of the lines doesn't really make a difference. But I'm quite sure that the target audience of your project doesn't consist solely of such people.

FEMALE VOICE: But in principle it is possible.

Leo: You get the idea. It just gets worse and worse and worse.

Steve: Yeah. Anyway, highly recommended. If anyone listening to the podcast has never run across on YouTube "The Expert," it goes on. It's about eight or nine minutes long. And really, it's great. And as you said, Leo...

Leo: Twelve million people now have seen it, so...

Steve: It's fabulous.

Leo: It's really good.

Steve: Yeah. It is good.

Leo: And it's real. You know? It's genuine.

Steve: Oh, yeah. In fact, my buddy works for 3M, and he's outside technical support. And I shot the link to him yesterday; and I said, "Mark, I bet you have been in these meetings." And he responded this morning by email, and he said, "Oh, my god," he said, "that was so funny, and it is so familiar." Because he's teamed up with a sales gal, and so she's pushing stuff that she wants them to be able to do.

Leo: Right. Of course we can do that, yeah.

Steve: And he's like, uh, no, you know.

Leo: You can draw seven straight lines that are perpendicular to each other. Of course you can.

Steve: All mutually perpendicular.

Leo: All mutually - no problem.

Steve: Oh, yeah. Anyway, it goes like that. It's highly recommended to our listenership. So anyway, this is the problem that we have, and this is what you see going on in Congress, because we have similarly, you know, in that video, the client wants what they want.

Leo: Right.

Steve: They want seven perpendicular red lines.

Leo: Drawn with a green pen.

Steve: Yeah, exactly. And with a green pen. And that's what they're saying, oh, but Silicon Valley, you guys are geniuses. You invent all this stuff. Just invent what we want. You know? And, oh. And in fact, I've seen some dialogue among crypto people who are dreading - it might have been some posts over on EFF, where they're dreading that someone will say, "Oh, yeah, okay, we can do that." I mean, probably the law enforcement folks have some people, some contractor who said, "Yeah, there's really no reason that can't be done. They just really don't want to." Because, I mean, you know, not everybody agrees on anything. So I'm sure you can find some people.

But anyway, Mozilla. Some changes coming to Mozilla. The first is that they've announced they're going to start requiring their Firefox extensions to be signed. And that initially made me a little nervous because, I mean, we know what a barrier signing is to things. I mean, that's the traditional reason that people didn't use HTTPS, and back in

the old days SSL, now TLS, is that they didn't want to have to get a certificate and go through all that. And originally certificates were even more expensive than they are today, and not available for free. The good news is that isn't what this is. It's just that Firefox is feeling, with the increasing breadth of their users, and sort of just the maturation of the platform, they're feeling more pressure from malicious add-ons. And while they can blacklist them when they are discovered and reported, they decided they need to be more proactive.

I mean, this is the model that we're seeing across the industry, the idea of being more of a curator and trying to perform some pre-release tests. So that's all this is. Rather than extensions just sort of being posted somewhere, and you can download them and install them, over the course of time, so as to give legitimate extension authors the opportunity to adapt, Firefox will be moving towards, sort of incrementally, towards this enforcement.

So, for example, right now I'm at 30 or 40. I think I just went to 40 the other day. Yeah, 40.0.2. So as of Firefox 40, Firefox warns about signatures being missing, but does not enforce them. And if you want to see that, just look at the add-ons tab in Firefox, and it sort of makes them a different color. They're kind of a puke-y yellow color. And it says along the top something about Firefox is unable to verify the integrity of this extension or something like that.

Now, when I just looked today, a number of them that no one had bothered to get signed were now signed. For example, NoScript. The first time I looked, when this was first announced, NoScript was unsigned. Now it is signed. And this isn't anything that the author needs to pay for. They just need to submit their extension to Mozilla to let Mozilla basically bless it. Mozilla looks at it, and in the case of NoScript I'm sure they just said, "Okay, yeah, fine, you know, Giorgio, we're not going to even bother with digging deep because we know you're a good guy." But so what's happening is, going forward, we're going to have safer extensions.

With the next release, 41, Firefox will have a preference that allows for signature enforcement to be disabled. And in fact the preference is already there. If you look in, if you have 40, which is the current one, about:config, under it's called xpinstall.signatures.required, it's currently set to false. And so 41 will flip that over, setting it to true. And but you'll still be able to go in and change it back to false if, for example, you're dependent upon some custom extension, maybe a corporate extension.

What I'm seeing is that this doesn't look like it's going to be a problem because, like, I have HTML Tidy, which is one I've just had around forever, that sort of audits HTML so I can quickly find unmatched tags and so forth. And it's not been updated. So it may be that we lose some which have been abandoned, but we're still using. And as of 42, Firefox 42, it will no longer allow unsigned extensions to be installed with no override. So with a couple more updates, if my little HTML Tidy that I've been using for so long - I'm sure there's, like, alternatives. I just haven't bothered to look because that's the one I've already had. I don't think it's going to cause a problem for anybody.

Okay. So that's the first thing. The second is Mozilla has announced that they are - and this was on the 24th, or the 21st, in a blog posting, so four days ago, that they are going to be adopting, moving forward, Chrome's Extensions API. And this is generally regarded as a good thing. You know, Mozilla and Firefox made up their own Extensions API years ago. There's sort of several of them that have wacky names. I'm sure extension people know what they are. But Chrome's has come on very strong. Opera and Chrome are both using the same one. And apparently Microsoft has been making noises about supporting Chrome's, it's called WebExtensions is sort of the generic name. And Mozilla has not said it will be 100% binary compatible. But they're saying that it will largely be compatible

with the model used by Chrome and Opera, and probably Microsoft Edge.

And so, again, that's only good news because what it means is that extensions developers no longer need to maintain two completely different code bases that operate very differently because the fundamental models of the browsers are so different. Firefox is going to go to a split process model. I don't think a process per tab, at least I hope not because that'll put me out of business with the number of tabs I have open, whereas Chrome does do a process per tab in order to get 100% inter-tab isolation. But Firefox is breaking the rendering process from the browser sort of root process in order to just - so they're doing some re-architecting.

So anyway, we ought to see that happening. And it's great. I love the idea that we're moving towards sort of a unified set of standards. It's been good for HTML. It was good for CSS. It's been great for JavaScript, as JavaScript, you know, Microsoft sort of made up their own Jscript with their own ActiveX extension stuff. That sort of died off. So we're seeing a unification, which just makes everyone's job so much easier.

One little annoyance that is upfront is that there appear to be some things that Chrome's current extensions API doesn't allow, which Firefox's does. I've already run across that on an app, on an extension that I've been playing with, that is cross-platform, that actually we're going to be talking about next week, which is something called uBlock Origin, when we start talking about HTML firewalls. And the point is that there are more features, it's able to offer more features under Firefox's API than under Chrome's. And they're things we'd like to have.

So maybe Chrome's will evolve over time, or maybe we're just going to have to give up on those things as they get unified. But at least not soon because what'll happen is that Firefox will add support for the Chrome extensions, but be much slower to kill off support for the old API. So that's a much longer, slower migration over, like, the course of another year, like through 2016 I think is sort of the timeframe of that.

Okay. I did want to, you know, the whole - you guys have talked about Ashley Madison and the breach on, like, probably every one of your podcasts, Leo, over the last couple weeks. And as you said at the top of the show, there isn't anything strongly security related to...

Leo: Just that they did a crap job. And we talked last week about that silly trusted security award on the front page.

Steve: Yes. And we could be annoyed from a technology standpoint that they claimed to - they charged people extra to proactively remove them from the database, and didn't. There was no sign that there was anything behind their promise and their delivery of we will remove you from the database.

But the point that I, in reading some of the extensive coverage across the 'Net, the thing that struck home for me that I did want to mention, because this also ties into another interesting topic about the power of Google search results, is the intersection of life and Internet technology; that what we have been witness to over the life of this podcast and the 10 years before is that the Internet is becoming something that, to a greater and greater degree, is part of our lives, that we really are becoming increasingly dependent upon. Now we're seeing services like, famously, ber, where it is connectivity and 'Net-based, I mean, as a fundamental function of it. And there are people who are using ber now extensively, more and more, because it's now available in their area, and the service

works.

So relative to Ashley Madison, without commenting at all on what it is, the point was, the point I thought that was really interesting was that, sure, we read about breaches where people's email addresses, and maybe their credit card numbers and some worrisomely personally identifiable information...

Leo: Their security clearances.

Steve: Yes. I mean, and it's bad when it's your images of your fingerprints, as the Office of Personnel Management, you know, 11 million people got those lost. That's not good. But it really hits home for people when it's, you know, those are not arguably life-ruining events. And to the degree that people actually registered their lives and have lives that could truly be ruined by the loss, the disclosure of this information, this is a sort of a different scale of impact on people's real-world physical non-Internet. It's not about email. It's right in the - smack you right in the middle of your family.

And so I thought it was sort of interesting, an observation that security and privacy as issues do increase in importance. We don't have solutions for them because, exactly as you said, Leo, they have security awards on the front. And, you know, these websites that have emblems about, you know, we were scanned this morning, and so we're secure, you know, that's - everyone knows that's just complete nonsense.

Leo: It actually makes me think they're less secure.

Steve: Yeah.

Leo: Doesn't it? Like, oh, these guys think I'm an idiot.

Steve: Yeah.

Leo: Yeah.

Steve: Yeah. So, okay. So that was one thing, just sort of this notion that, yes, security and privacy, I mean, the problem is we don't have regulation. And we know how hard it is to get it right. But I guess my point, it's increasingly important that we get it right for high-value privacy issues, and no one could argue that, you know, this has really hurt a lot of people, the fact that this particular type of dating site lost their information.

The other thing that I wanted to mention that is about the same thing was a recent paper published by the Proceedings of the National Academy of Sciences, PNAS, which was titled "The Search Engine Manipulation Effect (SEME) and Its Possible Impact on the Outcomes of Elections." Just the first paragraph of this, their abstract, they said: "We present evidence from five experiments in two countries suggesting the power and robustness of the search engine manipulation effect. Specifically, we show that biased search rankings can shift the voting preferences of undecided voters by 20% or more; and, two, the shift can be much higher in some demographic groups; and, three, such

rankings can be masked so that people show no awareness that manipulation has occurred. Knowing the proportion of undecided voters in a population who have Internet access, along with the proportion of those voters who can be influenced using SEME (Search Engine Manipulation Effect), allows one to calculate the win margin below which SEME might be able to determine an election outcome."

And then this was taken one step further because - and the link I have in the show notes, there, Leo, is the election 2016, the 2016 U.S. presidential election Google search trends. And so somebody who was involved in this research, actually one of these two researchers said: "According to Google Trends, at this writing Donald Trump is currently trouncing all other candidates in search activity in 47 of 50 states." There is a state map on that page where you can see.

"Could this activity push him higher in search rankings, and could higher rankings in turn bring him more support? Most definitely, depending, that is, on how Google employees choose to adjust numeric weightings in the search algorithm. Google acknowledges adjusting the algorithm 600 times a year, but the process is secret; so what effect Mr. Trump's success will have on how he shows up in Google searches is presumably out of his hands."

And we, you know, there has been discussion before about one of sort of the self-fulfilling problems or aspects of Google searches is that, since Google comes to know who we are, Google biases the results we see based on what Google has determined are our interests. And so it can be that an individual who is searching for things gets results that are more pertinent to them, but sort of - I don't want to say the "dark side," but the questionable side of that is what effect that actually has on them moving forward. And so, again, this is life intersecting Internet technology, where to such an extent we depend upon sort of the neutrality of results that we get from search.

Leo: Well, we'll see.

Steve: Yeah.

Leo: We'll see. I mean, I think Google has - we've always said this. Google has a huge responsibility.

Steve: Yes, yes. And I think that's, I mean, that's the best...

Leo: And absent good competition, there's nothing to keep them honest.

Steve: So another topic that I haven't ever touched on because I want to stay away from politics, but something in the last week happened that was a security lesson. And that's this whole question of Democratic presidential candidate Hillary Rodham Clinton and this issue of her email. And I'm a spectator of politics. I know that many people are a spectator of sports. For me, I'm interested in sort of - I love watching this human drama. And I've been, for weeks now, or months maybe it's been, you know, sort of watching this. And I'm interested because I know all about email security. It's a topic we've discussed extensively here. But there were never any facts. There was never, you know, it was just people on both sides talking about something they know nothing about.

There were questions about whether the server was wiped, but we were never - no one ever said what that means, what the term meant, how it was wiped, and so forth.

But in the last week came the information that did cause me for the first time to do a little bit of drilling. And that was when something was said about the server's apparently lack of physical security. And what made the headlines was that we heard that the server was in a bathroom closet of some small mom-and-pop Internet firm in - I don't remember where. Was it Colorado? Somewhere like that.

And so according to former employees - and I thought, okay, maybe it was they copied the stuff off to another server, and the unused one was in storage or something. But apparently these were in operation. An employee said the servers were in a closet off the bathroom. And one was quoted, a former employee, saying I don't know how they run their operation now, but we literally had our server racks in the bathroom.

And so that sort of brought me up short because for me to get to my servers at Level 3 I have a coded inductive badge, which I have to present to a reader. My right hand is biometrically measured, so there's biometrics. I have to enter a PIN. So I've got something I have, something I am, and something I know. And that gets me into the front door. Then there's conspicuous cameras everywhere, and I have to use a combination lock in order to access my servers.

And so the point of this is that I could withhold all judgment for lack of sufficient information about the actual security of an email server which I knew nothing about because there are certainly a large number of listeners to this podcast who are capable of building, algorithmically, a data transfer secure email server. A bunch of us know how to do that. But you cannot have data security if you don't first have physical security. And so I'm depending upon all of those measures to keep physical access from GRC's server out of casual people's hands.

And unfortunately, if in fact there was, you know, while the Secretary of State's email was being actively transacted, it was in an environment whose physical security seems very questionable, then there's a chain of delegated responsibility. Who knows whose fault it was, how that happened. But really it's impossible to defend the idea that that could be secure without physical security because, as we know, you have to have that first and foremost.

Google, as you know, Leo, because I've heard you talking about it, has surprised us all by announcing a router, the OnHub router.

Leo: Oh, yeah. I forgot I should ask you about this, yeah. Good, I'm glad you...

Steve: Yeah. And so there is one best feature that it has, which I immediately perked up when I saw this in their FAQ. And that is that the FAQ reads: "OnHub also automatically updates without interrupting your WiFi connection."

Leo: Mmm, love that.

Steve: "So you'll always have the latest features and security upgrades." And you know, we've had an acronym for a long time, of course, TNO. We're going to have to come up with another one. And I haven't given it any thought except you've already heard me

beginning to rev myself up about it. And that is, if it's connected to the Internet, the entity responsible for it has to be able to update it.

Leo: Oh, I like that, yes.

Steve: You know, that's what we're learning. That's the lesson of StageFright and Android phones and iOS phones. You've heard me complaining that the phone has Internet service. Somebody is generating revenue from the contract you have. With that must come the responsibility to update that device. Any device on the Internet. We've seen it now. We're used to it now with our desktop and laptop computers. They get updated. And we're moving into the IOT, the Internet Of Things.

And so arguably OnHub will be the hub for some future plans that Google has. And I love it that it is able to acquire new features on the fly. I mean, they love it because it allows them the flexibility to adapt to an unknown future. We don't know what's going to happen in the future, what's going to succeed or not.

But so on one side there are features. On the other side there's security. If it's connected to the 'Net, the entity responsible for it needs to be able to upgrade it because what we are learning is we are unable to do perfect security. No one thinks there's a single problem in anything they ship, hopefully. But they are always found to be wrong. And so it just - we just have to adopt, as a policy, if it's connected to the 'Net, then the person responsible for it, the entity that created it, has to be able to fix it when problems are found. Not if, when.

Leo: I love that. Very true. Very true. And I thought you were going to say the feature you liked the most was the lack of WPS, but...

Steve: I do appreciate that.

Leo: And I also figured, I figured there's not going to be any buffer bloat issues with this. These guys are smart enough not to fall for that; right?

Steve: Right. I mean, I assume that it's going to be good security out of the box. And what I'm comfortable with, if any problems are found, they'll just get fixed. I mean, and this is a model, obviously, that Google is very familiar with thanks to Chrome, which just, you know, it's fixing itself all the time. It's just, you know, that doesn't bother its users. I mean, we have to credit them with being first to get that, that, I mean, they're ahead of Microsoft, even, because Chrome just, you know, it just did it. And it turns out, well, that's really what everyone wants. It's just don't have a problem. And if you do, fix it. You know, the old-timers here, you know, used to download these things and look at them and roll them around in their mouth for a while and then go, okay, I think I'll swallow this. But that's just not the way of the future.

Leo: Not the way.

Steve: So briefly, two shows. I heard you talking about them on TWiT. I will echo the

thoughts of everyone who went wild with you two days ago, Leo, about "Mr. Robot." It's worth your time.

Leo: I started watching it last night.

Steve: Good. It won't disappoint you. It's just - it's great television. And we've covered it on the podcast, even before it started, because we'll remember that they released the first episode a month before the official series began. And it was so well received that the series got renewed for a second season before they even started the official series rollout. So, yeah. Really, it never disappointed, and the finale is this coming Sunday. So it's great. And the other show that I have mentioned, and that I know you have been watching, or watched because we're done the first season, is "Humans."

Leo: Yeah, just started that also. There's too much good TV.

Steve: There is plenty of good TV. And for what it's worth, I give that also a five-star rating. It ended with a wonderful tease for where they're going to go in the future. And people should know, I was not sure for the first few episodes whether it was going to get off the ground because it is not in a hurry. It developed nicely. But, boy, it's an absolute top recommendation. So for those who don't know.

I wanted to take a minute and ask our listeners who've been with us for a long time to bear with me because I'm beginning to get questions, and I saw them in the mailbag today, from newer listeners who are complaining that I'm talking about SpinRite being good, but I never tell anybody what it does. And it's not the case that I never tell anyone what it does.

Leo: Not at all.

Steve: No. On several occasions I have. But as it was Jeff Needles, your guy, and we've talked about this on the podcast before, over New Years he told me that Security Now! is growing at a nice pace. And what that means is that, if I haven't spoken, if I haven't explicitly explained what it is in the last couple years, there's a lot of people who have been hearing about it but don't know what it is. So as quickly as I can - because, again, I don't want to drive people crazy - about 25 years ago the problem with hard drives was that they were not interleaved correctly. They weren't optimized.

And the interleave is something that has fallen by the wayside, that's no longer an issue. But computers back then, and controllers, were so slow that the drives were spinning faster than they could accept, that they could either read or write the data. So sectors were spaced apart and interleaved among others so that the computer had time in between sectors to get that data into its main memory. So, but it turns out that the interleave was often as bad as it could be. The next sector was too soon, which meant the computer wasn't ready at the beginning of the sector, so the disk had to go all the way around again for the next opportunity to start reading that sector from the beginning.

So what the lesson was, was that, if you could adjust the interleave, if you could space the sectors out just a little bit further, then you would increase the performance

dramatically. So I first created SpinRite as an interleave optimizer, to optimize that interleave. The problem was everybody back then was already using their hard drive, and hard drives cost, as you mentioned last week, I think it was, Leo, I think you talked about a $5,000 hard drive. That's what they cost us. They cost as much as the rest of the whole computer and all of its other components. You know, really expensive. So we didn't have thumb drives or any place to put our data. And oh, my lord, you'll remember backing up to floppies. It was just, I mean, it was awful.

So I had to do an in-place re-interleaving of hard drives, which meant I had to reformat one track at a time. And if I was going to reformat one track at a time, I had to get all the data off the track before formatting it because this was the old-style low-level format that actually did completely erase the data. So our listeners know who I am. I'm a perfectionist about this kind of thing. And if I'm going to get the data off the track, I am really doing to get the data off because, if I'm going to format that track, I will never have another chance. It has to be now.

So what I built into that very first version of SpinRite was the strongest data recovery technology I was capable of designing. And over successive versions of SpinRite, that was SpinRite 1, we went to 2, we went to 3.1, and then of course 4, 5, and 6. And that technology, as the resources in the system have gotten stronger, I've been able to expand that. And at one point SpinRite, I think it might have been with v5, I finally took out the re-interleaving code. It was still there in case somebody had a really old system and drive, but finally it was just wasting space and code.

And what happened was that people began using SpinRite, not to do re-interleaving, but for data recovery. Because it turns out that thing that I built into SpinRite only because I absolutely positively had to get every possible bit of data off the drive, turned out to be the best data recovery anyone had ever designed in the history of the computer industry. And that sounds like Donald Trump. But it is, you know, it's the truth. There's never been anyone who has developed the data recovery technology that SpinRite has, all the way to the point where SpinRite is able to guess and check the guesses of data it cannot recover.

And it's even able to do partial recovery of data in a sector because it turns out sometimes, although you'd like to have it all, sometimes, if all you can do is get all but 20 bits, and the drive just refuses to give you that last 20 bits, that can still be enough in order to allow the OS to boot so you can get your data off of it, or allow your database to mount so you can then get all of the data out of the database. One way or another, what SpinRite does, and it is still unmatched in the industry, if anything will get your data back, SpinRite will.

And so what you've been hearing me, those who have never heard this explanation, what I normally just do is share people's feedback. They send us testimonials saying, "Hey, yeah, I heard you talk about it on Security Now!. Thank you. It saved my butt." So that's the story. Basically it doesn't do any of the things it started out doing. But it turns out that people still do lose data.

And the reason it has a future is that it turns out it works on non-spinning media just as well as spinning media because the nature of the market pressure, we're always talking about market pressure one way or another it seems here, in this case the nature of the market pressure causes the manufacturers to really strain the technology so that some percentage of the data stored will not be recoverable. SpinRite typically is able to recover even that, just by doing all of its tricks and trying really hard one last time.

Leo: Yay, SpinRite! Woohoo! Geez Louise. Well, go ahead. Let's do some questions. You ready?

Steve: You betcha.

Leo: Just forgive me for the noise. This comes, our first question, from Jim Sanders, Irvine, California. [Noise] Sounds like they're wiping Hillary's email server up there. Actually, that's what he's asking about. With all the news about Hillary's server being wiped, a question for you, Mr. G: What's the best way to do that? Hillary would like to know; right? I know I can run a hard drive over a degaussing device multiple times. That's what we do at the radio station; right?

But if I want to "wipe" a server, presuming that that means cleaning off all the remnants of files I don't want others to see, but I'm happy to keep certain files which are okay for them to see - oh, boy. Then that complicates it. What would be the best way to do so? I was thinking some process that might include writing multiple times on all available space, overwriting any latent data that might be recoverable. Or perhaps there's a better way.

Steve: Okay. So if you want to get rid of everything, we have a favorite utility - until I release mine. And that's called Darik's Boot And Nuke, DBAN. It boots, as its name says, so that it has full access to your hard drives. And then it does a very good job of wiping. And in fact, I would say too good a job of wiping because there is this, what's become more of really only historical, this concern of latent data being recoverable. Once upon a time, when we were using MFM and RLL encoding, there may have been a chance of residual data being recovered from underneath a wipe. It's just beyond feasibility today. Anyone who understands what happens between the user's data going in and the actual magnetic fluxes which are written on the drive understands...

Leo: Now they're doing your roof.

Steve: I've got the garbage truck dumping its garbage. It happens every week around this time. Anyway, there was once some belief that it was necessary to overwrite it 34 times, which maybe toward the - there was a period of time when we only still had 20MB drives, yes, "meg" drives. And computers were fast, and they were interleaved at 1:1, that you could afford the time to overwrite a drive like that. Now you just can't. And it's really the case that you do not need to.

But the problem is this doesn't really answer Jim's question because that wipes the entire drive. I would respond to degaussing, though. I think it's better, if you have a power drill around, to drill a bunch of holes right through the drive. That lid is always flimsy, easy to punch through, and then just keep right on going. They may be glass platters. They may be aluminum, depending upon the age of the drive. But it's not that hard to destroy what's inside.

The problem with degaussing is that you need to get really in contact with the surface in order for that to be effective. I doubt that you could do a degaussing that would be truly effective, both through the metal of the rest of the bulk of the drive, and also at the distance that the platters are, especially the lower platters further away from the

degaussing.

So, okay. So all that said, it turns out there is a little-known command in Windows called "cipher." It's been there since XP. And it allows - it's got a whole bunch of options. And it allows you to - it's a command line utility. So you would open a command prompt. And you can type cipher, C-I-P-H-E-R, space /? to get a list of all of its commands. And I think it'll do something different if you don't give it any parameters. So do give it a /?. My point is there's a /w command which is a wipe of unused space for Windows. And since I doubt anybody is still on Windows 2000, only I and a number of other real curmudgeons are still using XP, but everybody else, Vista, Windows 7 and so forth, there's a cipher command. You can say, for example, /w, I think that's :c, and that would be Windows using the file system and wiping all of the free space.

So it's already there, built in. And I would argue that a wipe is certainly good enough for anyone. What I will do for my wiping utility is several passes of high-quality pseudorandom data that has essentially perfect forward secrecy, so that no one will ever be able to figure out what was written. That's the way to do it state of the art. And then of course use the advances that I'm developing for SpinRite 6.1 to make it the fastest wipe possible. And that, of course, we call Beyond Recall, which is pending till after I get 6.1 out the door.

Leo: You're a busy, busy man. Question 2 comes from Grinnell, Iowa. Alexander is becoming a fan of Google Contributor: I've enjoyed the recent conversation about ads and site revenue versus user respect. How do you feel projects like Google Contributor will factor into the new world order? This solves one of my primary issues with Patreon-style funding. For me, it hasn't been a matter of money; it's a matter of time and human resource. It's simply not feasible for me to sign up for a premium account at every site/podcast/service I frequent. I've often thought that some sort of micropayments system that automated that for me would be desirable.

Google recently released the beta of Google Contributor. and I've been putting five bucks a month in, and I see a lot fewer ads on a daily basis. Usually these are replaced with a clear static image that thanks me for contributing. Plus I can log into my contributor account and see what I have paid whom over a period of time. This system, to me, is desirable as sites I more frequently visit get a bigger cut of my contributor dollars, while I don't have to juggle accounts at dozens of sites.

Steve: And I wanted to come back to this. We've mentioned it a couple times. I picked up on it from you, Leo, on one of your other podcasts. It might have been This Week in Google.

Leo: TWiG, it was TWiG, yeah, yeah.

Steve: Oh, yeah. And I thought, hey, what's that? So I went over. I signed up. And I have to say, I feel the way Alexander does. I noticed, for example, when I actually went over to iMore, I see - what I chose was the concentric circles, sort of pastel circles. And I wasn't really noticing them at first. But iMore's site is covered with them now, which means I'm not receiving the ads, I'm paying iMore. And I've been doing some work in JavaScript for the SQRL project in the last few weeks, and so W3Schools.com. I noticed when I'm there I'm seeing both banner ad-shaped things and rectangular down the side, and those are the same.

And similarly, as Alexander found, when I go over to Contributor, it shows me, I looked this morning, 560 ads had been replaced. And it itemizes how much money, my money, has gone to each of those sites in return for not having an ad and everything that we know comes with them. So I sort of wanted to just, in our coverage of alternatives and solutions and the dilemma, this is an interesting alternative. So, and of course, as Alexander says, other ads are coming from places other than Google. But Google, after a purchase of DoubleClick and similar acquisitions, they're certainly a large source of ads on the Internet. I have no problem with this. I like this. I think, you know, it feels right.

Leo: Well, I hope it takes off. Todd at Patrick Air Force Base in - where is that? Patrick Air Force Base. Well, the chatroom will tell me - wanted to follow up on Sandboxie. I was investigating the use of Sandboxie around the same time you brought it up on the podcast, followed by the issue with Firefox and how Sandboxie couldn't handle it. I found on the Sandboxie FAQ under "How does Sandboxie protect me, technically?" the paragraph with links as follows: "It should be noted, however, that Sandboxie does not typically stop sandboxed programs from reading your sensitive data. However, by careful configuration of the ClosedFilePath and ClosedKeyPath settings, you can achieve this goal, as well."

I was wondering if these configurations would address the issue you raised, or would it render browser use unacceptable? I heard your comments about a VM machine, but how would that handle downloads to the host computer? I've only dabbled in VMs, and I just haven't had the time to teach myself. Thank you for your time.

Steve: So I did want to - I wanted to come back to this because I had talked about how in the context of Sandboxie the problem was we recently saw this Firefox problem which Mozilla quickly patched, to their credit, and as far as we know, nobody got affected. But that was this PDF rendering flaw which essentially allowed scripting, JavaScript, to break out of the same origin constraint and essentially acquire your local file system as its origin, when then allowed JavaScript to have free access to the file system and rummage around and exfiltrate files.

So my point was that here was a problem that Sandboxie wasn't fixing, which is what essentially scared me away into considering, okay, we need really industrial strength sandboxing a la a separate virtual machine solution. And of course I've talked about, after thinking about that some more, some of those problems. And we've got a couple people who responded to that issue coming up in this Q&A.

But I did want to follow-up, saying that the Sandboxie people themselves heard the podcast and sent me a note saying, hey, Steve, you can change the configuration to block writes out, as well as reads - or, I'm sorry, to block reads from the file system as well as writes. And of course what they do by default is prevent you from writing. And here we didn't want a problem with exfiltration.

So I did want to follow up. I didn't want to leave that hanging with people believing that Sandboxie couldn't be fixed. I'm still unsure about what direction I want to take. But the next few podcasts we're going to be talking about some of these things.

Leo: Okey-dokey. By the way, it's in Brevard County, Florida. That's where Fort Patrick is.

Steve: Ah.

Leo: Yes, thank you, chatroom. So Todd at - it's not a Fort. Patrick Air Force Base, thank you, chatroom. Also in Florida, Robert Osorio. He's in Lady Lake, Florida. He reminds us of another solution for safe browsing: Steve, you've already described an important strategy for dramatically increasing browsing safety in the past, which is to make your Windows user account a standard user instead of an administrator. As you yourself have reported, this mitigates the vast majority of exploits. I think it was, like, 90 percent; right? It works so well that I have my wife's PC set up that way. I also use this on many of my residential clients' PCs.

The reason you aren't using it yourself, I'm certain, is because you haven't yet moved away from Windows XP on your main workstation, and XP doesn't play nice with "limited" users, as they were called back then. It breaks many apps in XP to run as a limited user, and there is no way to dynamically elevate your rights.

Windows 7 and 8 and 10 handle this way better. Now known as "standard" users, when logged in as one you can still run processes that require administrative rights. For instance, if you need to install a program, or make a system settings change, or run an older program that needs to run as an admin, UAC kicks it; but, instead of just asking for confirmation, it prompts for the administrator's password. I run my own laptop that way. And like you, I use a lot of old apps that violate modern secure practices, but they run fine as a standard user, although some require an admin password to run each time. For that reason I would recommend keeping your password relatively short and easy to type. After all, this isn't a situation where you're trying to keep out an intruder or a hacker - unless you keep your server in the bathroom. You're just creating a barrier to prevent malicious code from running on its own without administrative approval. And, yeah, that's actually a good idea. And I didn't realize UAC has been modified to let you log in. I just right-click and run as admin.

Steve: Yes. We talked about this back when we were first covering Windows 7. And probably our listeners will remember, and you, too, Leo, that when you log into Windows 7, a pair of credentials are created. It creates both an admin credential and, as they're now calling it, a standard user credential that has reduced rights. And it's the fact that they have this pair that allows you to relatively seamlessly sort of upgrade your rights as you need it, you know, when you need to do that. So I've thanked Robert for reminding me of that. And he's right that, of course, that's the problem with XP.

I did want to mention a feature of XP, or a utility that we talked about. It's funny because I thought - I was looking for some updates on it. It's a utility called DropMyRights. And it's sort of the reverse of that. It's a utility that allows you to run programs, if you would normally have more privileges than you want, to run a program under restricted privileges. And if you google "drop my rights," our podcast where that was the topic of the podcast is the first link. Now, maybe, again, my Google results are not what other people's Google results are. But that's the first thing for me that came up was our own discussion of DropMyRights, way, way long ago.

And in fact I had forgotten about it and had stopped using it. So I'm going to experiment with that. The author was at Microsoft, and it's a tiny little utility. And I have, in fact, I think I have a link. I'm sourcing the original file from notes that come up when you google "drop my rights" on the eventuality that the file from Microsoft might go away. It was a Microsoft developer. It turns out that Windows already had all this support, I mean

Windows XP already had all this support for creating restricted credential tokens, essentially. And then you create a process under that restriction.

So anyway, I'll have a little bit more to say briefly next week because this was all as I was putting the podcast together. I found Robert's note, thought oh, my goodness, we haven't talked about that for a long time. And again, this is only of interest for XP people, so I won't take much time on that, except to let anybody know if it actually ended up being something that worked. But, yes, that is definitely a benefit of moving forward to Windows 7 is that they have continued to fix this so that normal users are running with much more restrictions. And how many times have we, when talking about a Windows vulnerability, have you heard me say, you know, "allows malicious code to run with the rights of the user." That's always the case. It's breaking out into your process, you know, to your login rights.

And one of the things we have sort of forgotten is that, although not perfect, an operating system is under a sandbox. It is really working hard to sandbox individual processes from one another. And that's, of course, why Google is running a process per tab, is by having separate processes they're taking advantage of the OS's inherent process sandboxing. So if we put a browser in Sandboxie in its own lowered rights process, boy, we've got a sandbox, the browser's own internal sandbox, in a sandbox, in the OS sandbox. And is it any surprise we can't get our files out of there?

Leo: There's a lot of sand, I'll tell you.

Steve: Surprised it still works. Can it even find the Internet from in there?

Leo: Todd in Vancouver, Canada worked out a way to use VM-based browsers conveniently and securely. I've got a lot of email about this, as I gather you have, too, yeah.

Steve: It's a huge, I mean, we hit a real focus for our listeners, yes.

Leo: Hi, Steve and Leo. I have been listening for years. When I download my weekly bolus of podcast data - wow - Security Now! is always my first listen. Good. Top of the bolus. I also enjoy SpinRite, which has saved my butt both reactively and proactively several times. At my company we use the VM approach for browsing securely. Specifically, we use VMs using VirtualBox as the hypervisor. VirtualBox isn't a hypervisor, but that's probably a technical distinction and unnecessary. We faced the same issues you mentioned, specifically, how to deal with file downloads and how to deal with clicking links. We found approaches that I think are worth sharing.

Regarding file sharing, we set the VM up as a file server, but limit the sharing to the Host-Only Network. This virtual network exists only within the hypervisor so that the host OS can reach the VM OS, but nothing on the physical network can do so. This is important to protect the VM filesharing. It's important that the VM OS be set up as the server, since it is an untrusted machine. I have seen many users set it up the other way around because VirtualBox provides a convenient tool called Shared Folders, which basically sets up the host OS as a file server. You don't want an untrusted OS, the VM, to be able to access the host, only the other way around.

Regarding link-clicking, we use a utility to handle this. I gather he's an IT guy. It installs as the default browser in the host machine. When you click on a link, this application gets invoked. It is configured through a series of regular expressions - wow - to determine how to forward URL handling to different VM-based browsers or to the host browser. For example, in our company, Intranet URLs get forwarded to the host OS browser. URLs to external but core services requiring authentication go to a "secure" VM, and all other URLs to the "insecure" VM. This means links get opened in different browsers in different VMs automatically when you click the link. This has the added very cool benefit that the trick of spam that contains look-alike URLs doesn't work anymore - as long as you've properly crafted your regular expressions. The regex engine will never mistake a zero for the letter "O" or the number one for "L," and so on.

The last thing, which you didn't mention, is cut-and-paste handling. VirtualBox has the option of host-to-VM, VM-to-host, and bidirectional clipboard handling. This can be a security-sensitive issue, too, especially since there can be a tendency to cut-and-paste passwords from a host OS password manager into a login page in a VM browser. Ideally an insecure VM should be set up with VM-to-host-only clipboard sharing, if any at all. Cheers. And all I can say is I'm really glad I don't work at that company.

Steve: Well, first of all, I thought it was interesting that whatever Todd's company is, they have gone to this level, apparently on a per workstation basis. I mean, they recognize what we know at this podcast, which is unfortunately browsers have become a serious security threat. It's the way stuff gets into our network. We go out and get it. We solicit it by mistake. CryptoWall, CryptoLocker, these things get in by visiting a page that has malware. And I appreciated Todd sharing this. I thought his notion of the direction in which you shared was interesting.

I'm not quite as concerned as he when he talks about the direction of the file server because normally it's the VM itself which is sort of offering the server surfaces, rather than a server outside on the hosting platform. So he's right, it's on the host side of the host VM border. But it's normally managed by the VM supervisor itself, rather than by the operating system. Still, it's a nice point. And I did not know, because I haven't ever taken a close look - I'm sort of a VMware person - I've never taken a close look at VirtualBox. But this notion of the directionality of the clipboard is useful, too.

So, you know, many people - I've gotten as much email as you have, Leo, I imagine, about this. I mean, as I said, the mailbag was overflowing with stuff - are really interested in the issue of how do we lock down our browsers. So I appreciated Todd sharing his company's solution.

Leo: I figure it's got to be military or government; right? I mean, that is - or a bank. Maybe it's a bank. Robert Tomlinson in Austin, Texas suggests we take another look - and this is one I was going to bring up, but I'm glad he mentioned it - at the Qubes OS: Steve, I've been a longtime listener to Security Now!, and I'm a happy SpinRite owner. Recently you and Leo discussed your quest to secure - not me, just Steve - to secure browsing - I just surf like a normal person - perhaps in a VM, but without all the heavy lifting of needing a whole underlying operating system to support along with it, just being able to browse the internet safely.

I think there is a solution, you did mention it a few years ago on another episode, Qubes - Q-U-B-E-S - OS. This operating system secures individual apps in their own VM, each app in their own VM. And since your first mention of it, it's come a long way and now supports Windows app VMs without requiring the entire Windows operating system be installed in each VM space. I'd love to hear what your thoughts on this project are now and get a revisit podcast discussing it. Thanks for what you do. I anxiously look forward to my Wednesday commute every week so I can listen to Security Now!.

Steve: So we talked about this effort, as Robert remembers, and you remember, Leo, in the context of Joanna Rutkowska, I think is the way her name is pronounced. I wrote it down. But when I switched screens here, I lost my little note because I wanted to get her name exactly right. It was Joanna Rutkowska, I think. And she has - boy, and I also wrote down the name of her company. It's Tiny Things or something [Invisible Things Lab]. Anyway, she is a serious security researcher who developed this whole notion of really looking carefully at interapp isolation. And one of the things that - and I did go back, prompted by Robert, to take a closer look at it.

One of the things that they talk about is that it's possible to do this isolation with individual Windows applications without needing to boot an entire instance of Windows per VM instance. So I'm really interested to know how that happens because that's definitely interesting. The problem is that this falls in the category of, I don't know, I guess I would call them more burdensome to be practical solutions. That is, to actually implement this, you don't boot Windows any longer. You boot this Qubes OS, which is based on a flavor of Linux, and then it runs a hypervisor that then creates VM instances, and you go from there. So, wow. It's like, yeah, that would work. But if we thought that the previous question was overkill, or, like, went to great lengths to protect surfing security, this is more so because, I mean, this requires that you scrap your entire existing installation and start over.

Leo: Her blog is The Invisible Things.

Steve: That's it, Invisible Things.

Leo: And it's Rutkowska.

Steve: Rutkowska, thank you. So I'm looking for something, I mean, the word I'm trying not to say is "practical." But that's the right word. You know, something that, you know, as Leo's reaction was, "Yikes," you know, that seems like a lot to go through, all these VMs and link tracking and forwarding and everything. I'm trying to come up with something, you know, Script. You know, a solution that gives us what we want, but doesn't otherwise require that we start from scratch.

But I certainly acknowledge that rebuilding a system with Qubes, even though I haven't yet gone back and taken a close look at it, is certainly an option for, you know, if your need is so great that you need that kind of interprocess isolation, certainly Joanna has done that. And the project is moving along. It's like 3.0, I think. Or I think they're now getting ready to release v3.0, and it's been evolving over time.

Leo: And I'm sure she'd appreciate it if we say that she pronounces her name Root-kov-ska.

Steve: Root-kos-ka.

Leo: Root-kov. It's in her blog. She actually, I guess, gets...

Steve: Oh, no kidding, as if there's a "V" in there.

Leo: Yeah, Rutkowska. Joanna Rutkowska. And she works on Qubes OS.

Steve: Yeah.

Leo: So she's the one to talk of, to refer to in this context, for sure. Gabe does tech support for non-techies. He's got a solution: Steve, love all you do. I do tech support for non-techies for a living. I get malware-infected machines every day. For normal users, the only thing that works is to get a Mac plus Chrome plus uBlock Origin. Sandboxie's too complicated. No matter how well you set it up, they manage to disable it. He's like a zookeeper. He works with the monkeys. I've found even setting up a Hackintosh with the occasional maintenance cost when Apple does big updates is cheaper for people than having a Windows box and getting it cleaned when it gets infected.

As for Chrome OS - which is what I would recommend - it's simply not practical for non-techies. I don't know why he says that. I've got a 12-year-old using Chrome OS. I see first-graders using Chrome OS. If Google would get their act together, they'd make a Chrome OS VM that you'd run in Windows instead of Chrome. But of course, they don't make that. Cheers. Isn't Chrome the same, I mean, don't they?

Oh, by the way. P.S.: When setting up a Windows 10 machine, always, always disconnect the Internet. You save many steps, and it auto creates a local account without having to go through complicated hoops. So I guess it would have to, wouldn't it.

Steve: Okay. So, yes. Again, we are going to talk next week about uBlock Origin.

Leo: Oh, good.

Steve: Because I've been using it now for about a month.

Leo: Is it Mac only?

Steve: Nope, it's cross-platform - Safari, Chrome, Firefox, Mac, Windows, everywhere.

And Linux. Because it's a browser extension. And the way I will describe it in much more detail next week is one of the other things, Leo, I know you'll remember, is I was one of the early people, not only did I discover the first spyware and create OptOut, the first antispyware utility, but I was also one of the very early people pushing software firewalls. I understood the danger. That's what ShieldsUP! was all about was that people's computers were not protected. You need protection. And of course I ended up choosing ZoneAlarm at the time as my favorite solution because Billy had his act together. They had the technology nailed. And it was lightweight. It was the right solution.

What I think we need today is a web firewall. We need sort of, now, I don't mean everybody. I don't mean casual users. But the kind of person who would have once curated the settings on a software firewall, I think the right solution is for us to do that with a web firewall, essentially an HTML firewall. And that's really, it turns out, what uBlock Origin is. It's underdocumented. It's got an interesting history that I will describe. And I'm going to talk about the features that it supports. And unless something happens in the meantime, that's my planned topic for next week.

Leo: Ooh, I'm kind of liking this.

Steve: Yeah.

Leo: It's an ad blocker, in effect, because you can't phone home. You can't call out.

Steve: Correct. And it allows per-site customization - deep customization, if you want deep customization. It's very popular. Something like 1.3 million downloads by Chrome users under Chrome, and similar over on Firefox. So we're going to give it full coverage because I think that's the kind of tool for our audience.

Leo: Boy, it couldn't be easier. Geez.

Steve: Yes. It is very easy. And to get us back some control over what's happening. But the thing that really triggered me was his P.S.: First of all, I love that he does tech support for a living. He's dealing with actual users all the time and understands that, you know, this is what his solution is: Mac, Chrome, and uBlock Origin. That's the solution. And I'll explain why that's probably true next week. But Mark Thompson and I [dropout] yesterday, and he had set up a Windows 10 machine. And it turns out Windows 10 setup takes two entirely different paths if it has a connection to the Internet or it doesn't.

And I've not pursued it in any depth, but I wanted mostly for all of our listeners to hear that. For anyone setting up a Windows 10 machine, experiment with not letting it talk to the Internet when you set it up. Apparently it's an entirely different experience. It just behaves. And again, I don't have the details. But Gabe has just said that. "P.S.: When setting up a Windows 10 machine, ALWAYS [all caps] disconnect the Internet. You save MANY [all caps] steps, and it auto creates a local account without having to go through complicated hoops." And Mark Thompson, who you and I both know and is a deeply respected techie, he said exactly the same thing, which is why, when I saw that this morning, it's like, oh, I've seen that before. That's exactly what Mark said.

Leo: I'm impressed with how easy this is to use, and how transparent. I am, in fact, on a Mac with Chrome.

Steve: Yup.

Leo: Wow.

Steve: Cross-platform, cross-browser. And if you go to, in customization, you can go to a tab with a large variety of different curated lists of domains, which you're able to add or subtract to and from it, and it offers really good protection.

Leo: By the way, no requests on Rutkowska's page.

Steve: Yup.

Leo: Toggle strict site blocking. Toggle the blocking of all pop-ups. Toggle cosmetic filtering. Wow. So this is an ad blocker. But what it's really - it doesn't say an ad blocker. What it's blocking is third-party requests; right?

Steve: It's actually - I would call it, it's fair to call it an HTML firewall.

Leo: Yeah, that's a good [crosstalk].

Steve: Or an HTML filter. Filter's a little bit of a term, I mean, we understand what "filter" means, but firewall is a better, it's a stronger statement because essentially this is parsing the HTML, and with sets of rules that an advanced user can drill down. You can create regular expression exceptions and all kinds of things. There's a ton of power. But it also doesn't hit you with it. As you said, Leo, it just works.

Leo: Well, and I'm pleased to say our site works pretty well under it. It doesn't display ads, you know, but it doesn't...

Steve: Well, and see, that's the other thing. I think what's going to end up happening is we will see sites that start saying, if you want our content, you must disable your ad blocker. And the beauty is, Leo, that big power...

Leo: Easy to do, yeah.

Steve: That big power symbol, I don't know if it says - see, Chrome shows pop-ups, but Firefox doesn't. But that is a per-site disable. So if you disable it on TWiT.tv, go away and come back, it's sticky. And so the idea would be, this sort of solves that problem of going

to a site and either needing to or wanting to accept its ads and/or trust it to deliver ads. This makes that easy. This gives us control in a very simple way.

Leo: This is incredible.

Steve: Yet its default is to protect us.

Leo: Incredible. And very fast, by the way. It does not - and as a result of blocking all those inbound requests, it's speeding up things considerably.

Steve: I know. That's the other thing. In fact, from my standpoint, there are several - there's Adblock Plus. There's uBlock Origin. And there's a uBlock. And those projects got forked. One of the things that was compelling for me was that it is extremely lightweight. A guy named Raymond Hill, I want to say, I don't remember now if he's in the U.K. or Canada. I remember he wasn't around here. But he accepts no donations because he wants to feel no obligation to anyone. He feels like a little bit of a maverick by nature. But it is, as you mentioned, incredibly fast. It is the lightest weight in terms of minimal additional overhead for what it does of any of the filtering tools.

Leo: Well, let's not spoil your show for next week.

Steve: Yup.

Leo: I can't wait to hear more. And I will run it on my Chrome browser for the next week and see...

Steve: Great, great. You can provide some more feedback.

Leo: Yeah. Wow. I never heard of that. Brady in Idaho suggests an opt-out using Adblock: Steve, in the past I've heard you wishing for an Adblock you could turn on for select pages. I just wanted to make sure you were aware of Adblock's Blacklist feature. That's exactly what it does, and it's what I use, says Brady. Just enable Blacklist, and if a site's ads go too far, I click "enable on this domain," and I don't see it again. That is a good way to do it. I like that.

Steve: Right. And I believe it's possible to flip the sense of uBlock Origin around so that you could - so that by default it would be its shields are down, and then you go, okay, this is crazy. Because as a matter of fact, this has all been - you and I have been discussing this for a few weeks. I was watching you on TWiT. The camera wasn't switched or the feed wasn't switched to your laptop, but you were trying to show something, and it was covered with banner ads.

Leo: All the time.

Steve: They were, like, even covering up the controls.

Leo: Yeah.

Steve: And you weren't able to, like, you know...

Leo: It's very frustrating, yeah.

Steve: They had to sit there and close all these things.

Leo: Yeah. I was trying to show a YouTube video, and it had pop-ups all over the darn thing. And I don't know if this would work in that context because that's something - maybe it is.

Steve: Might be. My guess is it would do the job.

Leo: Wow.

Steve: Yeah.

Leo: Wow. Well, I don't, you know, I don't normally run adblockers, kind of on principle, during the shows. People say, well, gosh, we're seeing all these ads for other companies and stuff. And I say, well, we're showing a guy's website, and this is what he wants his website to look like. So far be it from me to change it.

John Crowther in Derby - oh, I'm sorry, "Darby" - U.K. comments on recent audio quality issues: Steve, this week I've listened to the three most recent episodes of Security Now!, and I'm sure I am not imagining things. It appears the quality of your audio is at times awful. I'm referring to your "down the line" voice, not Leo's voice, not the jingles. Jingles? I'm certain it's just your audio part of the recording. At times you suffer from what I call a "springy" voice and at other times your voice is clipped and jittery, with words or parts of words missing. Surely I can't be the only person who's noticed this, but if I'm the first to moan, I can provide you with timestamps of examples, if you're interested. Is there anything you can do about this? It's starting to get irritating. And I'm sure I've not noticed so many audio quality issues like this ever before. Best regards, John. Well, we know what's wrong.

Steve: Yes. I lost my T1s. And a number of people have commented. So I just wanted to make sure people understood. This was nothing I had any control over. For the first 10 years of the podcast one of the reasons, I mean, I loved that I had T1s. One of the reasons was it was unshared bandwidth, that is, the T1s were commercial-grade connections with bandwidth, and not a lot of bandwidth, 3Mb is all I had, actually 3.08 because they were 1.54Mb each. But the point was they were mine. They were dedicated. They went to an AT&T datacenter where then they were put directly on the 'Net.

Unfortunately, a few months ago they were discontinued. Essentially, it became impossible for me to get them. I was paying $466 a month. They were $233 each because that's the going rate for T1s. And I was happy to do it because I loved the fact that my bandwidth for this podcast was absolutely rock solid. And I worried when I switched to a consumer grade, you know, consumer bandwidth, a cable modem, despite the fact that I cranked it up to the maximum speed offered, it's not about speed. It's about lost packets. That "spring" that you hear is the audio compressor trying to fill in, trying to do the best job it can to put something into a lost packet. And you hear it with cell phone conversations, too. It's sort of the same technology.

So for what it's worth, there's nothing we can do. I'm sorry that it's not - that the bandwidth quality, the audio quality is not as good as it has been for the last 10 years. But now I'm among the rest of everybody with the cable modem bandwidth.

Leo: What's weird is you're actually having more problems than most of our hosts do. And that's what I can't really figure out. But your system, you don't remember this, but early on we also had problems like this, even with your T1s. Your system is complex. It's not just like a pipe coming in from the wall. Right?

Steve: Yeah. But there's nothing going on. The fact that we have large periods of time when there's nothing wrong demonstrates it's nothing happening here. I mean, I don't do anything while I'm doing the podcast. Nothing, there's no bandwidth usage here. Everything is super quiet. So it's upstream of me. It's completely out of my control.

Leo: You're far worse than anybody else on the network. That's the only thing that puzzles me.

Steve: Yeah.

Leo: And you're a business-class connection; right?

Steve: Yeah.

Leo: It's so weird. I think Comcast, you know, I see - I also have Comcast at home, business.

Steve: I'm Cox.

Leo: Oh, you're Cox. Okay, never mind, then. You see, it's not a bandwidth issue, it's a jitter issue.

Steve: Yeah. It's packets. If you drop the packet, or they get delayed, then there's, you know, the audio cannot go through.

Leo: Yeah. We had jitter problems before with you. You don't remember those days?

Steve: I remember we must have because I wrote, I started to write a solution for it. Remember the idea was going to be that it would always guarantee a final audio result that was perfect, even if packets were lost in real-time, the result would be perfect.

Leo: You were going to have a ring buffer, yeah.

Steve: And I - Speex, I was using the Speex compressor.

Leo: And that's why we didn't use it.

Steve: And you did the A-B comparison, and you were able to tell the difference. But, yeah. But then, of course, video, we switched to video and so forth.

Leo: The jitter was, like, 99, I remember, in the old days. It was really bad. And that led to packet loss. I'm wondering - hmm. You have - the modem you got from Cox, is it rented from them? Or is it yours?

Steve: No, no, it's a top-of-line Motorola DOCSIS 3.0, state of the art.

Leo: Did you buy one?

Steve: Yup.

Leo: Yeah. And the routers are the same as you were using, I presume.

Steve: Yup.

Leo: It's too bad. Oh, well.

Steve: Again, I mean, it can't be here because if you get a run of perfectly good audio, and we do, I mean, I'm hearing the same thing coming back the line, then it means it's not here because there's just nothing, you know, it would be - it wouldn't be intermittent. This is outside of my connection. It's between here and the Internet, unfortunately.

Leo: Yeah. Yeah, it's pretty consistent. It's throughout the show. And of course today you dropped entirely. You froze for five seconds.

Steve: Yeah.

Leo: I mean, it's not always that bad. But I hear - it isn't consistent like every 30 seconds. But I hear fairly consistent dropout throughout the show. And I don't know what to do. It definitely happened when you switched to Cox.

Steve: Yeah.

Leo: One of the things, it could be an interaction between how - we've had this problem - between - what are we using, John? Well, you know what, we'll do this off the air. Because we're still in the show. But don't hang up right away. Thank you, Steve. Everybody should go to GRC.com. That's where you can find SpinRite, the world's best hard drive maintenance and recovery utility, as well as all the great stuff Steve gives away, including this show and full transcriptions and show notes and everything.

GRC.com. Your feedback is welcome at GRC.com/feedback. But he's also on Twitter, @SGgrc, and has become quite the Twitter user of late. So don't hesitate. You can go to our site, TWiT.tv/sn, and subscribe there, or get downloads of audio plus video. We do video, so you can see Steve's smiling face, at TWiT.tv/sn, and of course wherever you get your podcasts. It's also on YouTube and places like that. Thank you ever so much, Mr. G. Next week...

Steve: My pleasure, my friend.

Leo: ...we're going to find out more about this really interesting solution, uBlock Origin.

Steve: An HTML web firewall that's just the perfect profile for our users who are interested in exerting some control over their experience of the web.

Leo: Thank you, Steve. We'll see you next time on Security Now!.

Steve: Thanks.

Leo: Bye-bye.