Listener Feedback #138

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-342.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-342-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here, and we have a lot to talk about, including the missing coffee episode, why third-party cookies are bad - or are they? - and a brand new product from Yubico. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 342, recorded February 29, 2012: Your questions, Steve's answers, #138. It's time for Security Now!

Steve Gibson: You forgot where you were for a moment.

Leo: No, no, I just wanted to give it a long Security Now!. There he is, Mr. Third-Party Cookie, Steve Gibson.

Steve: Boy, did we stir up a hornet's nest last time.

Leo: Oh, I bet we did. We're doing a Q&A, so I know I'm going to hear it. I've been hearing it all week.

Steve: I know, I know.

Leo: And I think rightly so. I just think it's a good conversation to have, and so we

will have it.

Steve: Yes. This was the wrong place to try to tell people that tracking wasn't a bad thing. That didn't go over very well.

Leo: We've been inveighing against third - you have been inveighing against third-party cookies as long as there's been such a concept. And I still have the question, what's so bad about third-party cookies, but I'm sure you will explain in the Q&A.

Steve: Well, I actually have an innovation of my own that I've never seen anywhere, and I'll share it in the Miscellanea section at the top of our show today, just something that solves the problem. And I don't know why no one has done it. But maybe by speaking it, there'll be somebody with some influence listening. I know that we do have influencers listening, so that would be great.

Leo: From your mouth to the influencers' ears. But we also have Q&A. I only have one little commercial. I can do it any time. So why don't you just get into the news of the week? How about that?

Steve: Okay. I like that idea. So - if I can find it. Where did it go?

Leo: I have it. You want me - I have all your notes.

Steve: Here I am, waiting for you to get ready, and I'm not.

Leo: Let's start with DNT, Do Not Track. Actually, that was the biggest news that I was really pleased that came out of the week's news was the Obama administration threw its support behind Do Not Track.

Steve: Well, and more importantly, I think, Google did. So we now have Firefox, IE, Safari, and Opera, and Google has announced they're going to support it, too.

Leo: And they have now some governmental clout behind it, so using it will have some meaning, which is equally important.

Steve: Well, yes. Everyone says, oh, well, but all that does is it just puts the DNT header in your query. It's like, yes. But again, let's not make the - what is it? Let's not make perfect the enemy of...

Leo: The good.

Steve: ...the good. And this is good. This is to allow people to express their desire. That's where we start. And we'll certainly be moving forward from there. It's the creep factor, I think. It's the sense that companies will exploit, if they can. And also I think it's people - you're in a position, Leo, of really understanding this very well. So you're like, well, yeah, but - you understand the implications. A lot of people who don't understand it, don't know the limitations of what could be done, and so they're a little more frightened by it.

Leo: Well, and now here's a question. Will Do Not Track be turned on by default on all these browsers? Or will we have to know to turn it on?

Steve: Really good question.

Leo: I bet not.

Steve: My guess is it'll be there, but not enabled.

Leo: So now we'll have to educate people.

Steve: Yup. I have the technology on my site, which I implemented back when I was doing the cookie project. And the technology, it sounds trivial because lots of people do it, but they all do it client-side. Mine was a server-side include of a little banner to alert visitors of things. And it's general purpose. And the idea would be, if somebody was poking around my site, and I noticed that they had third-party cookies enabled, I'd just give them a little notice and say, oh, by the way, you've got third-party cookies enabled. And if you don't know what that is, click, and I'll tell you all about it. And the idea was...

Leo: You did something like this before, I thought.

Steve: Oh, I've had it for years. I just haven't finished the darn thing. So it's on my short list. But I tend to get distracted by portable dog killers and ridiculous other projects. So, yes. So as you said, the administration is stepping up the issue of privacy. And what's interesting also is that the big online advertising group have said they're going to comply.

Leo: Now, they have to; right? It's a public relations thing; isn't it.

Steve: Yeah, it's having some traction.

Leo: I do fear for the future of the free and open and unpaywalled Internet at some point. We can balance privacy and monetization, I guess.

Steve: Yeah, and I would love to have some numbers, or for someone to provide us with

some numbers that was definitive about to what degree it really matters who we are because, remember, when you go to a site that's serving ads, that site gets revenue from advertisers by displaying the advertiser's ads, and extra revenue, of course, if you click on those ads.

Leo: Not necessarily, by the way. We don't charge for clicks, for instance.

Steve: No kidding. Okay, so...

Leo: And we get to charge more if we can tell advertisers more demographics of the people who are visiting. Or, and I think this is built into the price, the advertisers themselves know, because that's really what this is all about is DoubleClick, let's say, knowing who's seeing the ad and valuing that more. You see what I'm saying?

Steve: Yes. But we don't have any quantification of that. Because my point is that third-party cookies are the way that additional information is aggregated, but not the way the site presenting the ad is identified. That's the referrer header that goes along with the request for the ad. So DoubleClick, for example, knows that somebody went to TWiT.tv and is looking at one of their ads. Now, the question is, do you actually get more money if they know something about the user? Now, remember that they're putting ads on your site, so they know about in general your site's demographics.

Leo: Very general. Very general. In fact, we don't collect that information, so they don't.

Steve: Well, you don't. But that's of course what they're...

Leo: Well, but if they are keeping third-party cookies, they might. So that's the problem is we don't know, as a seller of ads, we don't know how that price is determined. We can set our own price, but whether it's worth that is not - only DoubleClick knows, or whoever buys the ads. We actually don't use DoubleClick or anything like that. We don't do that kind of thing.

Steve: I saw something years ago that said the whole concept of profiling had really not succeeded, although I think that's obsolete because I think Google has raised the bar because, when you put in a search query, right then and there you're saying this is what I care about. This is what I'm searching for. Well, that's why Google is Google. The size and the success they are is that their response page is able right there to show you ads that you essentially just asked for.

Leo: Right.

Steve: So it's not profiling you, it's looking at your query and saying, oh, we've got some advertisers who are paying for the use of the keywords that this person is just searching on. That kind of thing, that absolutely makes sense to me.

Leo: All I could liken it to is adblock. There are lots of our listeners block ads.

Steve: And this is very different. Yes, adblocking...

Leo: No, I understand, but I'm just - let me explain what I'm thinking, the analogy here. Such a small number use adblock that we can ignore it. We can safely ignore it. But obviously the person who uses adblock, and many people feel it's their right to use adblock, will not see ads and is not in fact contributing to - by the way, one of the reasons we monetize by reading you ads in the middle of the show is because we completely avoid this issue. You could skip it, but it's not automatic. But let's say we were making money, as we plan to, by the way. One of the ways - I told you we're developing The Tech Guy site. It's a $100,000 project. That is paid for by the banner ads that will be on The Tech Guy site. So this is a relevant issue to me. So if you use adblock, you're not paying your fair share of attention. But fortunately, very few people use it, so it's de minimis. If the default on the Internet suddenly became adblock, we would have to do business very differently.

Steve: Yes, everything would change.

Leo: Everything would change. So that's why the choice of defaults on third-party cookies might well be relevant. Again, we don't have the information to know. And we should get that information.

Steve: Would be nice to know. I know that Mark Thompson has a good friend whose site is SnapFiles, a really nice, high-quality, file-downloading site. And it's entirely run by - it's entirely financed by ads. The guy is making himself a fantastic living. And because the site has done so well, it's the only thing he has to do. And so he's able to spend a disproportionate amount of time actually looking at the freeware that he's posting, himself evaluating it, writing things up. Essentially, it absolutely closes the loop. The fact that he's displaying ads, advertisers are paying him for that display, the site is popular enough and so forth, I mean, it really does work. And so the question is, does it matter if the people who go there and see the ads are also being tracked?

Leo: Well, let's say he makes more money if they are than he doesn't. Who has the right to say how much money he makes? I think that there's this perception about that, oh, well, he shouldn't have to make that much, he shouldn't need that much money, he should just make what he can make without tracking. But that's not a decision for - that's a conversation we all need to have.

Steve: I always hail from a technical standpoint. And tracking is a mistake. This is a glitch. This was never supposed to happen. This is not why cookies were created, to allow advertisers to track people. Because the idea of a cookie is for you to have a relationship with the site you're visiting and allow it to maintain some state with you because you otherwise have a stateless relationship. And it was when this concept of a third party hosting content on the first party's site, when that happened, the advertisers realized, oh, hey, we're having cookies happening here.

And then, when that central repository of ads began being hosted on all these sites on the Internet, then that glued people together. It was that DoubleClick was a large advertiser who was providing third-party content across the Internet. Suddenly they were getting back cookies from people that they've given them to on other sites in the future. So that's where this whole concept of profiling users came from. I mean, this is not what cookies were for. And that's of course why browsers have allowed you to turn them off is it was recognized, wait a minute, there's a privacy concern here. So...

Leo: I don't think there's a privacy concern at all, but okay. I think that that's kind of just the default, hey, it's a privacy issue.

Steve: I would argue that, if someone goes to a Philharmonic website because they're planning to be in New York the next week, and browses around, but clicks nothing, then gets a phone call from the Philharmonics telemarketing company saying, hi, we noticed you were just over on the website and wanted to make sure you knew of some special offers we have coming up, I would consider that a privacy concern.

Leo: It is. But the concern is not from cookies. That didn't happen because of cookies. That happened because he gave his phone number to somebody who then revealed it to a marketing company. They were able to perhaps get that phone number from somebody other than the Philharmonic site. That's not clear. But the privacy issue is not the cookie. The privacy issue is somebody gave that phone number to the marketer; right?

Steve: The privacy issue is aggregation.

Leo: Yeah, but at some point he gave his phone number to somebody; yes? Somebody he trusted. They didn't get the phone number out of thin air. They got it from somebody he gave it to. Is that not right?

Steve: Correct, on some other website somewhere.

Leo: That's the privacy violation. Not the cookies.

Steve: So if he had third-party cookies disabled, they would not have been able to call him. That's my point.

Leo: Well, they wouldn't have known he went to the Philharmonic site perhaps, unless it was the Philharmonic that gave them the phone number, that he gave the phone number to the Philharmonic. The issue is that he gave the phone number to somebody other than the Philharmonic, visits the Philharmonic site, and then the Philharmonic's able to call him because that other site gave the phone number up. That's the privacy flaw. Otherwise, the only thing that's passed around is he was at this site.

Steve: People don't want to be tracked on the Internet. They don't. They don't want to be tracked.

Leo: Well, I think that they have to rethink that because they also want a lot of free stuff on the Internet. And I think that there's a real risk here that what - I understand what they're saying. I understand what you're saying. But I also think you need to understand that you're getting - I think people don't make the connection to all the free stuff they're getting on the Internet and advertising. That's how it's paid for.

Steve: Right. And I am 100...

Leo: You are risking undercutting all the free content you get on the Internet. You get a ton of it. Most of it.

Steve: We don't know that, though. I 100 percent agree with you about adblocking. Adblocking says I want the page scraped of ads and see the content. And I completely agree with you there. But that's completely different from tracking, from knowing who I am on two different sites that I visit, and having a third-party be able to glue that together and establish a profile.

Leo: Well, again, to me it seems like the profile only contains information that these other sites give up. And if they're giving up your personal information, your quarrel is with those sites. Yes?

Steve: Certainly that's not good, but that gets...

Leo: Yeah. So somebody - I gave my phone number to somebody, and somebody gave that phone number to DoubleClick, and then DoubleClick put the two together. But I really need to be mad at the person who gave my phone number to DoubleClick. Yes?

Steve: That's one part, Leo. Part 2 is, if I'm somewhere else, I don't want to be known as that same person who lost his information somewhere else.

Leo: You may or may not. And I agree that...

Steve: The cross-domain tracking, that's...

Leo: I agree that you should have the chance to turn that off. I'm not saying you shouldn't have the option to turn that off. However, I'm saying if it becomes a widespread option on the Internet, Do Not Track, it becomes the default option on the Internet, there may be surprising consequences. You'll see a lot more paywalls,

let me put it that way.

Steve: I think it's a little bit like what we experienced when I first discovered the first spyware and coined that term. And that is, users were furious that this had been done without their knowledge...

Leo: That I agree with.

Steve: ...and consent. And this is consent-free tracking. No one asks people if they want to consent to this. And so that's certainly a factor here, too.

Leo: Well, you go on the Internet. Do you not believe that you are tracked on the Internet? Do you not think that every site you go to keeps track of your IP address? Do you not think your IP addresses and all the sites visited are preserved by the Internet service provider? I mean, you're going on the Internet. Of course you're being tracked.

Steve: Okay.

Leo: And you're using free services. I mean, I agree, I completely understand, and I agree that you should have the right to turn this off. But I think that - and I think you have the right to use adblockers, by the way. I don't think adblockers should be banned. But I just think that you also should consider the fact that some of this is paying for the stuff that you enjoy so much on the Internet, including my content.

Steve: And it would be nice if...

Leo: And I don't think people want me to start charging you for content.

Steve: It would be nice if we knew whether tracking mattered there, and/or to what degree. We just don't know.

Leo: Well, it doesn't matter if somebody doesn't give my phone number to DoubleClick. It does matter a lot if they aggregate a lot, my credit cards and phone numbers, and then that information is shared all around. That of course matters. But that to me is the flaw, not the tracking cookie.

Steve: What people - it's the gluing it together. What people in the industry have said is that we would be stunned if we knew how much data was being collected about us. And so it's like...

Leo: All the time, by walking around, yes. I agree. Not just the Internet. We know this. That's how marketers work. That's how Kmart - I'm sorry, Target. You read that Target article. That's how it works.

Steve: And all of the little supermarket...

Leo: That's how they work. It's like saying I want a supermarket card - and by the way, Dvorak does this - I want a supermarket card, I want the deals of the supermarket card, just don't track me. It's a little bit more of an explicit relationship. It's obvious, I think, when you buy a supermarket card, that you're giving them all that information, including your phone number, everything you put on that form, matched to the products. That information is being shared with all the companies. I think you know that; right? But you do it because you get a deal on the food. So that's a much more explicit transaction. I just think we need to understand these implicit transactions that are happening on the web, that there is a reason for them. It's not people trying to see what you're up to. I mean, it really isn't. They don't care what you're up to.

Steve: Well, the theory is that, if you are profiled, then the advertisers can deliver more relevant ads, and that therefore the impressions are more valuable to you and to the advertiser.

Leo: Right. That's right. We don't disagree on this. And believe it or not, I think people don't completely - there's this kneejerk reaction, cookies are bad. You have to admit there's this kneejerk reaction, cookies are bad, and third-party cookies are worse. And I know lots of people delete their cookies routinely. And I think it's because they don't really understand what's going on here.

Steve: Well, and Firefox has an option of flush third-party cookies on termination.

Leo: Every browser does. Anyway, I don't care. Go ahead. It's much longer of a conversation about this than it deserved. It apologize.

Steve: Okay. So there's another controversy. And that is that Google, Microsoft, and Netflix are attempting to add DRM to HTML5. And the Mozilla people who are looking at this evolving HTML5 spec have been quoted as saying they think this is unethical. And Ars Technica reported on it, and the W3C has a spec and the working notes. And one of the things that they have asked is can an open source browser do DRM? To which the Netflix representative said, well, it's been the case in the past that open source software has included closed source modules, and/or maybe some hardware somewhere would actually be doing the decryption.

I was curious about where this thing went because, I mean, we know it's impossible. I mean, we absolutely know you cannot actually protect content. But Netflix's - Netflix - yeah, I said that right, Netflix's position is they would like - their content providers require them to protect the content that they're offering. And so traditionally they've used third-party tools or Flash or Silverlight or something that did offer that kind of

protection.

Well, as we're seeing Flash ebbing from the 'Net, and HTML5 is now becoming the solution, that and JavaScript, to glue together the capabilities, they're saying, hey, we'd like to be able to deliver protected content just to the browser without needing a third-party plug-in in order to provide that protection. So standing back from it, as we often said, if you're going to display unencrypted video on a person's screen, it can't be protected. I mean, the technology doesn't exist. It has to be decrypted in order to show it.

I have less of a kneejerk to this after really looking into it than the Mozilla people do, and I think they're sort of taking the ivory tower position because, in fact, what Microsoft, Google, and Netflix are asking for are just some hooks. They're asking to expand the media-handling API, the media-handling features, to provide the hooks that would allow JavaScript to request decryption keys and provide those to a decryption module. And those things are just sort of black boxes, not defined. They just want to get that into the spec.

And so my sense is, well, yeah, I can understand this. As a person who sort of really doesn't like having big, bloated plug-ins added to my browser, and loves the idea of things being kept simple, I like the idea of, if I want to watch something from Netflix, that I don't have to have a plug-in in order to do it, to provide the DRM as an add-on, but the browser can be enhanced with that at some point in the future. So what do you think?

Leo: Yeah. Oh, yeah. I mean, look. You're not going to get movies without DRM.

Steve: Right.

Leo: So good luck. One of the reasons Netflix wasn't on Android for a long time is because they didn't feel they could protect the movies.

Steve: Right.

Leo: We could argue about whether DRM's necessary. I don't think it is. But it's just the fact of the matter, if you don't support DRM, you're not going to get movies.

Steve: Right. So big news from the Eleventh Circuit Court of Appeals. And happy news for us. I want to read just the beginning of the appeal and then the decision because it actually refers to TrueCrypt. And I think exactly what was said and how it was said will be of great interest to our listeners. The document reads - this is an appeal of a judgment of civil contempt. "On April 7, 2011, John Doe was served with a subpoena duces tecum" - some legal term, a subpoena - "requiring him to appear before a Northern District of Florida grand jury and produce the unencrypted contents located on the hard drives of Doe's laptop computers and five external hard drives. Doe informed the United States attorney for the Northern District of Florida that, when he appeared before the grand jury, he would invoke his Fifth Amendment privilege against self-incrimination and refuse to comply with the subpoena.

"Because the government considered Doe's compliance with the subpoena necessary to the public interest, the attorney general exercised his authority under Title 18 USC, authorized the U.S. attorney to the apply to the district court, pursuant to Title 18, for an order that would grant Doe immunity and require him to respond to the subpoena. On April 19, the U.S. attorney and Doe appeared before the district court. The U.S. attorney requested that the court grant Doe immunity limited to 'the use of Doe's act of production of the encrypted contents of the hard drive.' That is, Doe's immunity would not extend to the government's derivative use of contents of the drives as evidence against him in a criminal prosecution."

So basically what they were saying was we will give you immunity for unencrypting your drive, but then whatever we do with the contents is still ours to pursue independently. So again he refuses. So this ends up being appealed, and the Eleventh Circuit just found that the Fifth Amendment right against self-incrimination does protect us against being forced to decrypt hard drive contents. They wrote:

"We hold that the act of Doe's decryption and production of the contents of the hard drives would sufficiently implicate the Fifth Amendment privilege. We reach this holding by concluding that (1) Doe's decryption and production of the contents of the drives would be testimonial, not merely a physical act; and (2) the explicit and implicit factual communications associated with the decryption and production are not foregone conclusions.

"First, the decryption and production of the hard drives would require the use of the contents of Doe's mind...." This is what you and I were talking about before a couple weeks ago when this was still circulating around, and we weren't sure how it was going to come out, Leo. So it would "use the contents of Doe's mind and could not be fairly characterized as a physical act that would be nontestimonial in nature. We conclude that the decryption and production would be tantamount to testimony by Doe of his knowledge of the existence and location of potentially incriminating files; of his possession, control, and access to the encrypted portions of the drives; and of his capability to decrypt the files.

"We are unpersuaded by the Government's derivation of the key/combination [lock] analogy in arguing that Doe's production of the unencrypted files would be nothing more than a physical nontestimonial transfer." Where in the past it's been argued that the government could have the combination to a safe, for example, or the keys. "The Government attempts to avoid the analogy by arguing that it does not seek the combination or the key, but rather the contents. This argument badly misses the mark." And then they quote some case law.

And they finally talk about TrueCrypt: "To be fair, the Government has shown that the combined storage space of the drives could contain files that number well into the millions. And the Government has also shown that the drives are encrypted. The Government has not shown, however, that the drives actually contain any files, nor has it shown which of the estimated twenty million files the drives are capable of holding may prove useful. The Government has emphasized at every stage of the proceedings in this case that the forensic analysis showed random characters. But random characters are not files; because the TrueCrypt program displays random characters if there are files and if there is empty space, we simply do not know what, if anything, was hidden based on the facts before us. It is not enough for the Government to argue that the encrypted drives are capable of storing vast amounts of data, some of which may be incriminating.

"In short, the Government physically possesses the media devices, but it does not know what, if anything, is held on the encrypted drives. Along the same lines, we are not

persuaded by the suggestion that simply because the devices were encrypted necessarily means that Doe was trying to hide something. Just as a vault is capable of storing mountains of incriminating documents, that alone does not mean that it does contain incriminating documents, or anything at all."

So that's the upshot of that, which is good news for people who want the right to keep their stuff private and have the Fifth Amendment protect them from self-incrimination.

Leo: Yeah, we covered a little bit of this on This Week in Law, if people want to know more. It's very interesting.

Steve: Oh, cool.

Leo: And TWiT, as well, yeah.

Steve: So Yubico yesterday released a new form factor for their famous YubiKey, and it's very cool. They call it the "Nano," and it is essentially just the USB plug portion that would go into and disappear into the USB slot with a little bit of an arced metal contact and a light that can be seen. So the idea is that this is - it's sort of like a semi-portable Trusted Platform Module. We've talked about...

Leo: You're going to lose this, though.

Steve: Yeah. Well, exactly.

Leo: It's to be put in something else; right?

Steve: Well, yeah. It's to be put into - the idea would be, everyone's running around with laptops now. This would be a...

Leo: Oh, just leave it in there.

Steve: Exactly. You stick it in, and it just lives in the USB slot. And people who have seen this have said, well, wait a minute, that's not really multifactor because...

Leo: It's permanent.

Steve: It's permanent. But it authenticates the device, which is really nice. I mean, for example, I came up with that very fancy system using cookies, secure HTTPS-flagged cookies, in order to allow my employees to access the GRC data when they were roaming. And the idea was they had to have their laptop at home with their home IP that GRC recognizes. Then when they use the laptop to open a protected page on GRC, it says, hey, you don't have a permission cookie, do you want to get one? So then I give

them this permission cookie, which then identifies the laptop. Then when they're subsequently roaming, and it sees a query come in carrying that encrypted protected cookie, then along with them identifying themselves using the Perfect Paper Password system - and in fact I think maybe that's why I designed the Perfect Paper Password system was for this purpose, because I wanted something like a one-time token, and I also wanted their access to be locked to their laptop so that they wouldn't be tempted to just log in from some friend's computer because there's just - I'm not convinced there's a safe way to do that. So I wish there was some simple, clean way of authenticating that laptop. And this provides that.

You would still have them provide - first of all, it'll identify itself. And you of course touch the contact in order to send a one-time password for further identification. And then for multifactor you'd still pop up and prompt them for a password. And only if all of that works as a whole do you then say, oh, welcome, and give them what they want.

The other cool thing is that this is the v2.2 of the YubiKey technology, which has added a full FIPS standard/challenge response mode, which is to say it can have a secret which it never reveals. And rather than merely sending out a one-time password every time you touch it, it can be challenged purely through software and generate a response. So that, for example, you can use it to authenticate a submission. You put together a bunch of information, and you press the Submit button. Well, that information is run through the YubiKey. It generates essentially a signature which it sends off with the information, and that's able to authenticate what was submitted.

And in fact the well-known free and open source Password Safe technology now supports the YubiKey, not just this little guy, but all of the v2.2 YubiKeys using this challenge/response system. So you stick your YubiKey in your computer, and it now identifies itself while you're using Password Safe, which is over on SourceForge, a nice multiplatform solution. And Password Safe is able essentially to ping the YubiKey whenever it wants to and as many times as it wants to to verify that it's there and that you are who you are saying you are. And then you're also able to configure it so that you have to touch the little YubiKey contact if you want to set it up that way.

So anyway, I'm really happy. They're moving forward on this authentication technology. And I think the idea of being able to sneak one of these, just sort of slip it into a USB socket to provide some authentication for the laptop itself, is cool. And I might still use a YubiKey in a different USB socket in order to authenticate me separate from that. So, very neat.

And, boy, I hesitate to open this up again, but I did want to say that one of the solutions that I had always had and been thinking about for third-party cookies is just to associate the third-party cookies with the first-party site. I had this in my notes, so I wanted to mention it, and that is...

Leo: It's true. If Google hadn't used DoubleClick as the originating site, but had instead used Google as the originating site, none of this would have happened.

Steve: Correct. Correct.

Leo: So it's that second originating site that confuses the issue.

Steve: Yes. And my thought was, people have said, well, if you disable third-party cookies, some add-ins on sites no longer work. And that is the case. Facebook apps are running on your Facebook page. They're inherently a third-party app, and they need a cookie in order to provide the same, not tracking, but the same sort of stateful connection to the user that the first-party cookie provides to websites. Because I'm certainly with you, Leo, cookies are not themselves inherently evil. Nothing works if you disable all cookies. I mean, almost nothing anymore.

Leo: Well, the only thing we disagree on is how dangerous third-party cookies are.

Steve: Correct, correct. Yes. And so one thought is that, if third-party cookies were tethered to the first-party site where you were when you got them, then everybody's happy. Then you're able to accept third-party cookies. Third parties are able to associate with you, yet they're not able to track you to follow you as you go to other sites on the Internet. When you go to those sites, because your cookies would be in a jar essentially, and I coined the term "cookie jars" for this, then that other site would provide third-party content, and you'd have a third-party cookie for it. But it would not be the same third-party cookie because you were on a different first-party site. And that means that third parties are able to associate with you, no apps break, yet no one needs to worry about being tracked from site to site. So sort of it's a nice compromise.

I also wanted to say, as I was going through the mailbag, so many people said, hey, Steve, you can solve the Perfect Posthumous Passwords problem, which we've talked about, what happens if I die, how can I give access to people I care about. I was really impressed with how many people said LastPass, just use LastPass's one-time password. It'll generate some. You put them somewhere safe. And if anyone else ever needs to access your LastPass logins, they're able to do that using your first-party cookies, so I just wanted to - or using your one-time password. So I just wanted to shout out to everyone who also thought of that.

And a very brief Honor Harrington update. I thought I was off the hook when I finished Book 11, Leo.

Leo: Now, tell me the truth. Did you breathe a sigh of relief? Or were you sad?

Steve: Well...

Leo: When you say "off the hook," it sounds like you were glad.

Steve: Oh, my god, 11 books.

Leo: That's quite a commitment.

Steve: Yeah. It really did bog down sort of like two thirds of the way through this first 11. I mean, he's built a huge universe. You learn all this about all these people. David Weber, the author, clearly is interested in politics, so it's very much star system politics and bad actors and good actors. So it's like, okay. And, I mean, I was enjoying the

battles that just seemed very clean and interesting and fun. And there weren't a lot of battles there around two thirds through the series. However, I have to say that Book 12 has started off, and I'm immediately gripped by it, and not a shot has been fired. So I realize that I bought this series hook, line, and sinker. It's a major investment, as you say, in time. I'm ready to be done, but now I have to find out what's going to happen.

Leo: Now you know why I didn't start. Okay.

Steve: Ah, yes.

Leo: Now 12, a dozen. But more to come, you think? I mean, is this it?

Steve: Yes, yes, there's 13 and 14. 13 is supposed to be tomorrow, well, tomorrow's month, supposed to be in March sometime is No. 13. And then I think there are a couple more in the arc. And we are going to get movies, so people who just don't have any interest can just wait. It's going to be a few years, I'm sure. But you can see the movie.

Leo: So long as it's not Angelina Jolie's leg, I don't mind.

Steve: That was quite a leg.

Leo: Starring as Honor Harrington.

Steve: Oh, yeah. So Al sent me, from the U.K., just a short note reporting his success with SpinRite, which allowed his external hard drive to be imaged. He said, "Steve, thanks for the useful podcast that you do with Leo. I really love the show and methodical detail for each topic, right down into the ones and zeroes of computers and security. This is just my simple SpinRite story. I purchased it so that, if anything did go wrong, I'd have it to save me. I've experienced the gut-wrenching feeling of having a drive which is broken, and at least with SpinRite nearby I can feel calm that it can often fix the situation.

"I like using MS Flight Simulator to fly different planes and have it installed on an external USB Iomega hard drive. After a few minutes of playing, the drive becomes pretty hot, and I know that heat is not too good for a drive. Nevertheless, it seemed to keep working many times, so I thought that probably the heat may not be so bad. I don't know what the max temp is to run the drives so that they have a long life. Then one day I thought that I should make a backup image [audio dropout] because I have so [audio dropout] and in case the drive failed. I used Acronis TI, or Acronis, I guess, TI...

Leo: Acronis, yeah.

Steve: Acronis, to do this. But it told me it could not make the backup because there was a problem with the drive. I guessed that it might be damaged sectors, so I ran SpinRite over it, and it fixed one sector, and then the drive worked completely perfectly

again. I tried using a fan blowing air over the external drive while I played the sim, and it seems to keep it much cooler. Maybe that will give it a longer life. Thanks very much for your great product." So thank you, Al, for our report.

Leo: I just have this vision of him with his computer open and a big fan blowing on it.

Steve: And I have to say, Leo, in my own experience with hard drives, it's very surprising how much cooler drives will run with air.

Leo: Oh, yeah. Keep them cool.

Steve: It's not only keep them cool. But, I mean, I just had - just like sometimes I'll put a heat sink on the drive, like a heat sink with fins, and then blow the air across the fins. The drive almost feels like it's below room temperature, I mean like cold cold, cold to the touch. So by all means, I endorse the idea of air flow across drives. That just makes them happy.

Leo: Well, and a well-designed case will do that. It's a little tougher in a laptop, of course.

Steve: Well, and many external cases, I think, really don't...

Leo: Yeah, they're tight.

Steve: Yup. They want to make them small. And if they're small, you can't get much air through there.

Leo: Are you ready now, my friend, for some questions and answers?

Steve: Absolutely. We've got some good ones.

Leo: You've got the answers; I've got the questions, starting with Jean-Matthieu Bourgeot in Tarare, France, who also loves coffee, apparently: In the last episode of Security Now! you mentioned you and Leo had been speaking about coffee for 30 minutes before the show. Has this discussion been recorded anywhere? I would be greatly interested to listen to it. We've got to do a coffee podcast. The demand. I'm sure many other listeners would love to hear it, as well. Along those lines, why not do a special episode about the health benefits of drinking coffee, as you did for Vitamin D. Over the years you have been and continue to be a great inspiration to me in my work and a real plus in continuing my tech education. Thanks so much, Jean-Matthieu. Thank you, Jean, for listening.

Steve: So this caught my eye because I was very - I won't say I was surprised. But Leo, well, I asked you before we began recording, was our discussion somewhere captured? Because if it could just be stuck on YouTube somewhere, I know that it will get a chunk of our listenership interested in what it was you and I were talking about. So if anyone can find it, or if they have it or something, that would be great.

Leo: Well, that's a good question. If you go to Justin.tv, anytime we don't record a show, Justin.tv does. That's, by the way, a very handy thing, as it turns out, whenever I forget. Now I'm getting - look at that, I'm getting ads. How dare they. So if you go to Justin.tv and search for TWiT, we have a - you can watch our live broadcast, obviously. But what happens with Justin.tv is they record everything as it happens, and I think they do it in chunks. I'm not sure. How big are the chunks? They used to be two hours. I think they're longer. So you can go back to last week and get - I think we talked from 11:00 to 11:30 Pacific. So you should be able to on here, I'm told, I've never done this, but go back to our recording of seven days ago and catch it. Let me just, you know, let's - I don't know. I don't see it here. It should be here. They're days in length now? Ah, that's what they changed. They used to do it in chunks.

Steve: Wow.

Leo: But now they do it in bigger lengths. So one of our chatters has posted this link. Oh, boy.

Steve: Big link?

Leo: Let me click this. Well...

Steve: We need to bit.ly it.

Leo: It's Justin.tv/twit/b/309424801, if that helps in any way. And I guess they said if you go into this 38 minutes in - is that what they said? Something like that - you will find our conversation. So let's jump ahead. Yeah. No, wait a minute, that's the Tech Guy Show. So I don't know if this is - one of the things Google Analytics will tell you - I don't know. As you can see, everything is recorded. Everything is recorded that we do on this network. And I don't know, this looks like Saturday. So I'll have to go to a different link. I'll leave it to you.

Steve: So it's there somewhere. Maybe someone will find it and sort it out.

Leo: If you're in the chatroom, you've got it. 66:28 in, they said. All right, let me look. I'm clicking the link, and I'm going to go 66:28 in. 66:28 in. Let's see. Let's see if it worked. Oh, yeah, look.

[Begin clip]

Steve: ...focus point of the Internet. I mean, we all need DNS in order to resolve...

[End clip]

Leo: That sounds like the regular Security Now!. No, it's in there. No, they're wrong. I don't know why they said that, but that's wrong. So what you need to find is the raw - is not a repeat. Anyway, I leave it as an exercise to the viewer. And I apologize, but this is a security show, and we just really didn't think that you all wanted a show that was half an hour of coffee.

Steve: Well, actually you and I, we establish our Skype connection and then sort of chat a little as you're pushing buttons and getting monitors and cameras and things arranged.

Leo: Somebody needs to watch live, if you care about that stuff.

Steve: And we just sort of stumbled into a discussion, since I had invested all this recent time and energy in coming up with the perfect cup of coffee.

Leo: I can guarantee you, had we put that in the show, we would have gotten far more emails saying how dare you, it's a security show, stop talking about coffee.

Steve: Among the people who really expressed an interest, there was one grumbly person who said...

Leo: There's no reason to grumble.

Steve: ...oh, my god, don't. Don't, don't, don't.

Leo: We didn't put it in the show.

Steve: Okay.

Leo: And I'm not doing a special just for a half-hour of pre-show. My advice to everyone is watch live. That's all I'm going to say. We are working on it. I think this will be soon because I understand that one of the arguments against watching live is if you're not in a U.S. time zone, it's ridiculous. You can't watch live, you're from Australia, one of our viewers, and he's fast asleep when we're doing the show. But here's what we're going to do in response to that because this is going to kill a couple of birds with one stone. We are - I don't know why it's taking so long, but we're working on a way of just recording the whole day, from the moment I start to the moment the last show ends, and then flipping it, and the reruns are that day again and again. So you can watch, in your time zone, you just start watching whenever you want, you watch for eight hours, you'll get all new stuff, and then you

can go to bed again. How's that sound? Right now we don't do reruns - we do reruns of the edited shows. So you would have never seen that coffee stuff in the rerun.

Steve: So you'll take the eight hours and then duplicate it eight hours and eight hours to fill up a full 24.

Leo: Yes, exactly.

Steve: That's brilliant. That's wonderful.

Leo: That way, yeah, I thought about that a while ago, and it's just hard to implement. I guess we're having difficulty finding something that can record eight hours nonstop. I think we have something, but my engineers won't let me use it. They say the world will end.

Keith Rollin in Sunnyvale, California wonders, what's wrong with 2048 bits? Steve, on Security Now! 340 listener Craig indicated he thought that 1024-bit public keys should be secure enough. Perhaps I imagined this, but it seemed like he didn't like the idea of soon being moved over to 2048. Well, so what? What's wrong with 2048-bit keys? Why not use them, given they're more secure? Ever since I read "Cryptonomicon" I've used 4096-bit keys whenever I could. Am I doing myself a disservice by doing so? Should I strive for some balance between more security and fewer bits?

Thanks for a great podcast. I commute 15 seconds - this is not a typo - 15 seconds every day from my bed to my computer, and I couldn't do it without listening to Security Now!. It takes him a year to listen to one show, but - that's very funny. Thank you, Keith. It's a good point. Why not just use all the bits you can?

Steve: Yes. It's because they do not come without some cost, Leo. When we double the length from 1024 to 2048, we much more than double the computation required. It's a minimum of five and as much as 30 times more computationally burdensome to have 2048-bit asymmetric keys than 1024. So one of the things that's been sort of holding people back is already, as we know, people have been leery about the processing overhead of dealing with setting up SSL connections. It's historically the reason that websites bounced people into an HTTPS session only while they were logging in, and then bounced them back to a nonsecure session, the belief being that it was only during that period of time when they were actually providing their credentials that there was any need to protect their communications.

Now we know, of course, that the token that they were given went to create a session when they were logging on. Then that's going to be available in the clear. And that's what things like Firesheep were able to grab in order to impersonate people. So we know that we're moving more towards HTTPS all the time, SSL everywhere. I mean, that seems to be the future because it's going to solve these problems. But what's been making people in big data centers nervous is that means every single TCP connection is going to have to do this SSL negotiation. Google's been looking at it because they see this as slowing things down.

The good news is that SSL, as it's been more robustly implemented, allows the credential that is once negotiated to be reused. And so that goes a long way towards solving the problem. And next week I'm finally going to catch up with my commitments about shows that we're going to do and discuss the SPDY, the so-called "speedy" protocol. I keep getting people asking me, both through Twitter and in the mailbag, hey, Steve, when are you going to tell us about SPDY, which is Google's tweak to HTTP, specifically to address these sorts of issues. So we did the show on TCP and why TCP connections are expensive due to the need to throttle bandwidth, and that was sort of a preamble for being able to discuss SPDY and what Google has done in order to further improve the performance of the web. And I'm really happy with this R&D arm that Google has that's making this stuff go better.

But bottom line is 2048-bit keys are much more computationally intensive. On the other hand, the computing power in our machines is still going up exponentially. Now we're in multicores and multichip multicores and huge on-chip caches. So I don't think it's really going to be a problem. And it's certainly good to have the security.

Leo: Indeed. Let's see, next question. Question 4, is that right? Did I skip 3? Where's 3? Where did 3 go? Here's 3. David Jones. Oh, by the way, sad moment. Davy Jones passed away, the lead singer of the Monkees, at 66 today. He had a heart attack.

Steve: Wow, too young.

Leo: Too young. Way too young. But another David Jones, this one in Aurora, Illinois, and presumably still with us, says: Anonymous brings DNS down? So what? Steve, let's say someone does figure out a way to actually bring the root DNS servers down, as we talked about - this was last week's episode. I can't imagine that the major search engine databases don't have the IP of the sites they have indexed in their respective entries. Couldn't Google and Bing simply tweak a line of code to put the IP addresses in the links in the results page, instead of the domain name? Oh, that's interesting.

So if you can't remember where a site is, just search for it. And then click the link, and it would go to the IP address. Back in the old WebCrawler days there were lots of sites that didn't even have domain names. Then all you would need to know is Google's IP address. You would need that, wouldn't you. And you could still get to all of your favorite sites, as long as those sites don't change their IPs during the outage. Which I would think would be a pretty low priority project if there's a major attack on DNS going on. Let's not change the IP address today.

Non-web-based Internet services such as email or other apps might not have a way around these theoretical problems. I don't know enough about the guts of those services to know. Couldn't the rest of the planet's DNS servers just change a quick setting to ignore all DNS entry expirations until it's over? Am I missing something? That's one our most popular phrases, in almost all the email you get. "Am I missing something?" Thank you so much for the podcast. Loving them since Episode 1. Have a great day. Longtime listener and SpinRite evangelist, David Jones, Aurora, Illinois, the City of Lights. P.S.: How's the Vitamin D working for you? Any updates?

Steve: Okay. So technically David's right. But WebCrawler days? Come on. Yeah.

Leo: We've come a long way, baby.

Steve: Now, the problem is that a web page doesn't look like it used to look in the WebCrawler days, where it was just Times Roman text pouring down.

Leo: Static, right.

Steve: Exactly. We know that contemporary web pages are full of other URLs with domain names going back out and causing our browser to load all those resources. For example, many sites now just don't run without JavaScript, as we know. And the page, each individual web page, at the top of the web page typically, sometimes at the bottom, calls out the JavaScript that it needs from the server, or from a server, maybe even a different server. In fact, many times it is. You'll see JavaScript libraries being pulled, for example, from Google or other locations.

Leo: Oh, yeah. Oh, yeah, all the time.

Steve: A jQuery, for example.

Leo: Right.

Steve: Some people don't keep their own local copy, they just always go get the latest one. So the problem is you might put the IP address of the main site you're going to into your browser and bring up some skeleton of a page. But, boy, you wouldn't be using the Internet in that case.

Leo: Images are often from a different server. I think you do that, don't you, on your page, have a separate image server?

Steve: For a time I did, yes.

Leo: We have a separate MySQL server. So, I mean, in many cases, yeah, your page would break.

Steve: And media stuff is coming - I have media.grc.com is where all of the podcast audio is stored, and video stuff.

Leo: I suppose you could, in your bookmarks, you could, instead of having a URL in your bookmark, you could put an IP address. If you thought this was going to happen. Was it March 31st was the day?

Steve: March 31st, the day before April Fools Day, is when the claim was, although it's been disavowed by Anonymous folks, the claim that this was going to happen. And, I mean, also to David's point, there is no mechanism for suddenly making all DNS servers stop expiring their records. But that would work, too. But the problem is we have to have some assumptions, and we have to have some foundation. And the founding assumption is we have DNS. And the DNS system is diverse and spread across the globe and well wired up, and we're increasing the security of it all the time. Everyone takes the need for DNS very seriously. And it seems to me that focusing on keeping DNS up makes more sense than lots of little tricks to deal with what happens if it goes down.

So it certainly would be the case that an individual who was worried about this could start, could build their own DNS cache, start statically recording the IPs of all the DNS lookups their system does, and if there was ever a problem with all external DNS servers, just switch over to his own private cache of DNS names and IPs. That would work. Sort of like a big hosts file, which is exactly what that was used for once upon a time. So that's sort of a way of - on the other hand, you'd be sort of lonely on the Internet because you'd probably be there all by yourself.

Leo: Just be you.

Steve: Oh, and Vitamin D. I haven't been sick for the last several years, and I get email from people from time to time pointing out articles on Vitamin D. It's by some weird coincidence that, if I had to have picked one thing to recommend to our listeners, even today, although my research has continued and is continuing for many years since, if I had to choose one supplement, I would choose Vitamin D over everything else.

Leo: I'm getting mine out right now. I forgot to take it today.

Steve: And all of the research says - I will mention one thing, though, and that is that I finally saw a paper which showed the measured blood level of Vitamin D correlated with the amount of international units of Vitamin D taken by people per day. And so it was a scatter chart that had a series of vertical lines where the horizontal axis was the amount of Vitamin D being taken. So the reason these lines were vertical is that Vitamin D was being taken in increments, like 1,000 IU, 2,000 IU, 5,000 IU, 15,000, 20,000, so forth. So thus they were vertical lines.

But what really stood out in my mind was the degree of scatter of measured blood levels at any given daily dosage of Vitamin D. Some people really don't need any. Some people really need a lot. And it is a potent hormone. If you were to take 100,000 IU every day for a few months, that would damage you. You do not want to do that. But on the other hand, apparently if you drink endlessly water every day, that would damage you, too. So, I mean, anything in high excess will hurt you.

But the point is the only way to know where you are is, next time you get your routine physical, just ask your doctor to have your Vitamin D levels checked. The good news is, since we did the podcast, not in any way because of the podcast, but the world seems to be waking up to the importance of Vitamin D. So your asking your doctor is not going to have him go, "Huh?"

Leo: No, I was very pleased because after that thing that you did, I did ask my doctor, and he said, "Yes, of course." It was a simple, easy thing to do, and I was getting blood work done anyway. And he said, oh, yeah, yeah, we'll add that to the test. He didn't mind. It can't be a very expensive test. I mean, it's...

Steve: Not expensive. And what I learned was I was really, really low, even though my health has historically been very good. But, I mean, Leo, just except that I got bad food once, or actually twice in a restaurant in the last two years, I haven't been sick a day. And we did hear after the winter following that podcast from many of our listeners who said, wow, this is the first rainy winter season I ever went through that I didn't get sick. So it has strong immune benefits, and I can't recommend it enough. But you really need to test because I need to take a lot of it in order for my blood to be where I want it to be. Other people may not need any, or much less.

Leo: I got this test done March of last year. I should get another one done because, as you can see, I was at 29 where the standard range is 30 to 100. In other words, at the very low end of normal. And so I have been taking D since. And I agree with you, and it's purely anecdotal, but I have been much healthier since I took it. And when I forget for more than a few - I think you build up a level. But after a few weeks I start getting sick again. So, in fact, I just - I haven't taken it in a couple weeks, so I'm taking some now. Gotta take it. So, yeah, I mean, I was glad that you did that bit. And so consult your physician, of course. We're not doctors.

Steve: Yes, exactly.

Leo: We just play them on the radio. Question No. 4, Dax Mars, visiting Earth via Second Life. He says he never updates his computer: Mr. Steve, I've been listening and watching Security Now! since the start. Always good info, but I have a question. I never update Windows on my PC. Well, the one I use only for programming. I write shareware, you see. It's a Pentium 3 running XP Pro SP2. It's probably been four years since I did the updates as of that install. Yeah. I think there's a SP3 you're missing. It is on my wired LAN, but it is never used to go online except to allow the CD ripping software to get CD info once in a rare while. So should I be updating it?

My reason for not updating is to keep it stable as a programming PC without things changing all the time. It only has a 10GB drive, by the way. Of course that's more than enough space to code on. Of course, I'm not completely crazy, and my two main PCs are 100 percent up to date, run Avast!, SeaMonkey with noscript and adblocking, behind a Linksys WRT54G with DD-WRT firmware and one of your killer long random passwords on the WiFi. So should I update the Pentium 3 or let her be?

Steve: You know, I'd let her be. I think that it's been four years, you've never had a problem, you only go on the Internet in order to let your CD ripping software go and grab from the CD database, album name and track names. As long as you're not promiscuously web surfing - and also, by the way, you have SeaMonkey with script blocking. So you certainly want to treat that machine carefully because there's a lot of stuff we know is on the Internet somewhere that could get it if you clicked on a link. On the other hand, if you behave yourself, you have your other machines for that kind of work, I'd say, eh, I mean, I can understand the notion of leaving it alone.

One of the things that really annoys me about Windows is that every single one of these little updates, and they're often not so little, has the ability to be rolled back. Well, that requires that Windows stores all of this state of the machine before it applies each one. And although the directory is hidden under the main Windows directory, oh, my lord, after a few years, it is so full of stuff. And Microsoft just sort of, I mean, I don't know what they're going to do about it. I wish they would expire the older things after a while.

But it does annoy me the degree to which all of this just junk builds up in Windows, and it never goes anywhere. You have the ability, of course, to roll all the way back. But I don't think that is ever going to work for anybody because there are just so many dependencies and interdependencies among things. So it's one of the reasons setting up Windows from scratch is a good idea. I would say, as you suggest, Leo, just jumping to SP3 isn't going to destabilize anything, and you'd get a bunch of fixes there at once.

Leo: It's on the LAN, and other computers on the LAN are on the network. So is it really fair to say it's not on the network, just because he never points that computer's browser to a web page? It is on the Internet isn't it? It's behind a router, but it's on the Internet.

Steve: Oh, yeah, it is. And so he definitely has to behave himself. But it's a little bit like you and me, who both don't use third-party AV stuff. We're just very careful about where we go and what we do and, knock on wood, have gotten away with it so far.

Leo: I think it's different, Steve. I hate to...

Steve: Really.

Leo: ...argue with you twice in one episode.

Steve: No.

Leo: Because, all right, let's presume there's an unpatched exploit on XP. And that exploit can be - now, if course it would have to get through the router; right? So...

Steve: It would be him going to a web page.

Leo: There's nothing just kind of floating around that would go float through the router to his machine or could be on another...

Steve: Now, that's what's so nice. That's what's so nice about being behind a router and...

Leo: So the router's going to protect it.

Steve: And SP2 has its firewall turned on by default, too.

Leo: Right, right. So he would have to explicitly open himself to an exploit using a browser or some sort of surfing.

Steve: Correct.

Leo: All right. So, yeah, that makes sense. I don't disagree with you. I don't. I agree. It is a little different than AV because an antivirus doesn't protect - if you have an exploit on your system, it doesn't protect you against necessarily an exploit being taken advantage of.

Steve: Right. It tries to be watching your email come in and finding...

Leo: In other words, we'd be nuts not to patch it.

Steve: Correct.

Leo: It's one thing not to use an antivirus. It's another thing entirely not to patch your system if you are going online; right?

Steve: Right, right.

Leo: Omri Amirav-Drory was present at the Solve for X event that we've been talking so much about. In fact, Omri has been here in the studio, and we're going to have him at some point on one of our shows.

Steve: Oh, no kidding. Genetic compiling.

Leo: Yes. He's one of the speakers at the first-ever Google Solve for X event. His talk is online at YouTube. In fact, if you just search for "Solve for X" and "Omri," you'll find it. He met, he says, with the spray-on nano-particle antenna guy, Anthony Sutera, and his fiance during that event. He seems to be a very genuine geek. He brought samples of his material to demonstrate and showed several interesting electrical properties of the material to anyone interested. And, yes, he actually did have it in that spray can.

I am no chemist, so I can't verify his claims. But it's one hell of a complicated hoax if it is one. I'm a big fan of you and Leo, thought you and your listeners might appreciate another first-hand report since it's easy to wonder about something so new and surprising. Omri Amirav-Drory is a Ph.D. and founder and CEO of the Genome Compiler Corp., which is, interestingly, GCC is the acronym. He is going to join us on a triangulation to talk about genomics. It's fascinating what he's doing.

Steve: I'm glad he's going to because he did indicate in his note that he was available.

Leo: Yeah, we've been in touch, yeah.

Steve: Very good, very good. And I just wanted to say thank you. I appreciate having a report from someone who was there and who met Anthony and had a chance to spray an antenna on himself and so forth. So we'll just sort of see what happens with that. I don't have any stake in it one way or the other. I thought it was interesting, made a great presentation, certainly. And we'll just keep our eyes open, as we are for other things like supercapacitors.

Leo: I'm still waiting for that to happen. Jim Hartz, New Brunswick, New Jersey says he had his life made easier: I'm a huge fan of Security Now! and a SpinRite owner. I just wanted to thank you for something not related to either. I recently was doing a web search for how to disable UPnP on a Windows XP box, and the second result was for none other than GRC.com. Hey, nice SEO, Steve. Knowing the site, of course I clicked that link. Your freeware UnPlug n' Pray was quick and easy. Thank you for making my life a little easier and for all you do for your followers. Keep up the excellent work. That's a free program you offer. Does it work for all versions of Windows, Unplug n' Pray?

Steve: The ones that need it. The reason I pulled this, what caught my attention, was that I realized how long it has been since I've had to solve one of Microsoft's screw-ups.

Leo: That's true. Things have changed, haven't they.

Steve: That's my point, exactly, that I wanted to discuss a little bit, is the world really has changed. It's not like we don't have problems anymore, but they're different problems. And they've sort of moved. What I was doing for quite a while with DCOMbobulator, Unplug n' Pray, and, I mean, patchwork, one thing after another, was Microsoft would have some problem, and I'd quickly do a nice little lightweight piece of freeware written in assembly language, that no one could believe how small it was. It would just come up, they'd click the button, it would solve their problem. It would just fix that, whatever it was. And, I mean, that's one of the reasons I was on your early Screen Savers shows so many times. And this was happening all the time as I was producing these things. And that's stopped. I mean, that really has stopped. And I think it was them finally turning the firewall on in XP with SP2.

Leo: XP SP2, yeah.

Steve: And the fact that routers have become ubiquitous. I mean, once upon a time people actually plugged, I mean, I just shudder to think of actually plugging a computer directly into my cable modem, or into an ADSL connection. I mean, the idea of there not being this little separate island firewall appliance, which is essentially what a router gives us, as the first thing that's going to stop any unsolicited inbound traffic. And we do know that, even today, if you took an XP machine, like just XP the original "gold" build, and stuck it on the raw Internet, it just gets taken over by Code Red and Nimda and MSBlast

and everything.

Leo: Is that still the case, I wonder?

Steve: It's still out there on the Internet, yeah. It just immediately becomes taken over because there was no firewall, and there was no protection, and there was all, I mean, and DCOM things, there was like - Microsoft was running services that people didn't need, like Universal Plug & Play, which were vulnerable. There were mistakes in those services. And when you went on the Internet,, they were just wide open and exposed to exploitation, which is what made it such the days of the Wild West back then. But anyway, Jim's comment made me realize, huh, I'm not doing that anymore. I'm doing other cool things that I'm quite happy with. So, but yeah, I thought it was worth just sort of noting that the world has changed. That's behind us, thank goodness.

Leo: Such good news, yeah.

Steve: Yeah.

Leo: That's one of the reasons we can talk so much about privacy nowadays, frankly. Ed Zucker, in Long Beach, California - remember we used to talk about spam a lot, too. That's kind of gone away.

Steve: Yeah.

Leo: Ed Zucker - or not gone away, we've just found ways to deal with it, let's put it that way.

Steve: Right.

Leo: Ed Zucker in Long Beach, California, demonstrates that Chrome side tabs are well missed. We talked to you last week about the fact that Google had removed them. It was an experiment they turned off: Steve, you're not alone in your love for open web pages and tabs. The "bug" of the missing side tabs exploded in the chromium bugs management. And after 108 individual postings by people who were desperate to get side tabs back, Google was forced to close the topic to further posting. They didn't put the tab back, but they at least said you can't talk about it anymore.

Steve: We're not going to talk about it anymore.

Leo: We're done.

Steve: Yeah. I have a link there to this so-called "bug" report for Chromium.

Leo: It's not a bug, obviously, it's a complaint about a feature.

Steve: Yeah. It's like, hey, Chrome updated and my tabs are gone, my side tabs are gone. I want them back. And this thing goes on and on and on, 108 individual people saying, c'mon, figure this out, give them back to me. Now, the good news is the 109th posting was from a Chromium person. And what he said was that, sorry, they're not coming back, so we're going to shut down this complaint log because there's no point in continuing. But he referred to a technology that they were looking at which would enhance Chrome such that it no longer had a rigidly fixed Chrome, for lack of a better term. Remember that "Chrome" is the insider web designer/web browser jargon for the window dressing, all the so-called "chrome" of the browser is all of the stuff that's not the web page, the various buttons and menus and so forth.

And so what they're talking about is extending the extension, Chrome's extension API to allow customization by add-ons that would be powerful enough to allow tabs, side tabs. And that's the kind of thing we've seen in other browsers that have more permissive APIs that do allow more customization of the browser real estate. So that would be cool. That would mean that users who weren't tab addicts, as I and so many other people are, people who didn't need those wouldn't be burdened by having Chrome lugging that technology around. But those of us who absolutely organize our lives around having 57 open tabs would be able to add that to Chrome and then tab ourselves to death.

Leo: The developer post says this really - "The bug tracker is where the engineering team discusses bugs. This really isn't the appropriate place. We suggest you go to the chromium-discuss mailing list, a much more appropriate forum." And I have to say I think that that's probably the right way to handle that. And we should point out that, while Google releases Chrome, that Chromium is in fact not a Google project, it's an open source project that Chrome is based on. So this guy Pete Kasting, who posted this, and the people who made this decision, may or may not work for Google. It's not known. And I think having - it's a simple way to add it. Just add detachable surfaces or something like it.

Steve: Yeah. Isn't that great?

Leo: Then you can do it.

Steve: Yeah.

Leo: I mean, 108 sounds like a lot. But I'm sure there are close to 100 million users, so it's not - it's a small percent. Brian Voeller in Medford, Oregon, wonders about factors and prime collisions: Today I have a question on factors and an unrelated thought on prime number collisions, something we were talking about a couple of weeks back. You mentioned in Episode 340 that 512-bit and 768-bit numbers have been factored. So we're always going to move to larger numbers to stay ahead. That's why we're at much larger numbers.

When you say that they've been "factored," does this mean that someone has

decrypted an early test message after many months and years of brute-force computation, thus proving the math is correct? Or has someone computed all the products of all primes in the given keyspace and arranged them in something like a rainbow table? It seems that the former is not all that serious since all that brute-force work is completely useless against any other encrypted message, even if the time required will steadily decrease as computation power increases.

The real problem would be the latter, as decrypting any message would then be as easy as retrieving hashed passwords from an unsalted database with rainbow tables. The only catch I could see with this is that the tables would be so large that only a fraction could be stored in all the hard drives in existence. So just brute-forcing would be faster, perhaps, than searching. Can you explain a bit more about what you mean by saying they've "been factored"?

Steve: So it's really interesting. When RSA, the security company whose conference is ongoing right now, and the source of a lot of our crypto technology from its inventors, when they proposed an asymmetric crypto system based on the difficulty of factoring large numbers, which were composed of primes, they said, okay, we want to make sure these are as hard as we think. So they formally created a challenge which they hosted and gave, I think it was a $10,000 cash award for factoring the so-called "modulus," which is the composite of the two primes, breaking it back down into its primes. And the contest went on for quite a number of years. And it had a long list of increasingly long products of primes. And on their site, which is still available at the RSA Labs section of RSA.com, I think if you Googled something like "prime factorization challenge" you can probably get right to it.

And what happened was the shorter ones got factored, and the largest composite of two primes that was ever factored was the 768-bit composite. And I have the text from the discussion of that particular solution, to give people a sense for this. So this is 768 bits. "A six-institution research team has successfully factored the RSA-768 challenge number. While the RSA factoring challenge is no longer active, the factoring of RSA-768 represents a major milestone for the community. The factors were found on December 12, 2009 and reported shortly thereafter. The academic paper describing the work can be found at" - and then there's a link to a PDF. "The factors are" - and then it lists these two big, long, three lines of decimal digits is the two prime numbers. And it says, "The effort took almost 2,000 2.2GHz Opteron CPU years, according to the submitters, just short of three years of calendar time."

Leo: Oh, boy.

Steve: So in '09 it took three years for this six-institution team to crack a single 768-bit composite, a "modulus," as it's called, the composite of these two primes. No one has ever done it for 1024. And it's not like a quarter harder because 1024 is a quarter longer. It's exponentially harder as these numbers get larger. And so put that into perspective with 2048. I mean, it's just not happening. Which is why what we talked about a couple weeks ago, the discovery that the primes being used were not as random as we assumed they were, why that's a big deal. Because essentially the 2048-bit - in fact, this was mostly done with 1024. Even 1024-bit numbers are so hard to factor, nobody has ever - nobody has ever factored one, ever. But the idea that there's this shortcut where you don't have to factor them because you can use a Euclidean algorithm, if you happen to have two of these moduluses that have a common prime, it'll tell you almost immediately

that that's the case. That's like, whoops, that's not a good thing. So that was a real backdoor behind this.

But anyway, the idea was RSA had this challenge. Moduluses up to 768 were factored. No one has ever been able to factor anything bigger, which gives us a sense for the security. And that 768-bit one, it used all the cleverness these guys could come up with. The paper just makes your eyes cross, it's so detailed, and the technology that they used. And even so, 2,000 2.2GHz Opteron CPU years were required, and it took them three calendar years to crack one. And that was 768 bits. So no one's using that anymore. We're all at 1024. And we're headed to 2048. In fact, all of my new DigiCert keys are 2048 because that's the minimum length that you must use if you want to get an extended validation, an EV cert. You have to have 2048 as the minimum, which everyone expects is going to keep us safe for a long time.

Leo: So it's really being prudent. I mean, it would take three years for a 768-bit key.

Steve: Yes.

Leo: So it's just being prudent.

Steve: Yes, 2009 wasn't that long ago. We haven't gotten things that much faster. We still haven't had any breakthroughs in factoring. I mean, this is a problem that really smart people have scratched their heads about for a long time.

Leo: Question 9, Jacco Flenter in Holland wonders about salts and hashes stored in databases: Steve, I've been listening to Security Now! for quite a while, and I've got a question. I'm curious about hashes and salts. I'm looking to help a project which stores salts in the database together with the hash of the password. Now, is that a security issue? I understand that this would leak some information about how to generate the hash from the plaintext, should the hacker know what method was used, and I assume that that's a bad thing. But exactly how good or bad is this practice? Unique salts give also a bit more strength to the hashes stored in the database, since one rainbow table matching a salt does not corrupt the whole database of passwords. Thanks for Security Now!. Regards, Jacco. So he's saying there's a single salt, and it's stored with the database. Is that a mistake?

Steve: Well, yeah. He's saying that apparently they generate a random salt for each account. And so they're storing the salt and the hash of the account password.

Leo: I get it, together.

Steve: Together. And so, okay. What we know is that a salt is some typically binary data which is appended or prepended to the password, which are then together run through the hash to produce the result which is stored in the database. The normal way, the reason one's concerned about this, is the brute force, that is, you are guessing - if you had access to the database, that would give you the salt and the hash. If you knew what

the hashing algorithm was, you would then start putting guessed passwords through the process, trying to get the hash to come out. If you could do that, then you would - you don't know that you have the password because hashes can have collisions, as they're called. But you would know that you had a password which, when combined with a salt and hashed, gave you the result. And that would allow you to impersonate that user, to log in as them into that system.

If the salt was not available, what's significant is that the salt is binary, yet the passwords you're guessing are ASCII. And if the salt were not available, if it was stored somewhere else, or if it was encrypted or somehow better protected than stored right next to the hash, then it's very likely that no ASCII, no password characters could be put in that would create a collision with the hash, that it would give you the same hash because you would very likely have to put some binary in, in order to have it salted with the binary which is the salt in order to get the result.

Maybe I've made that little too confusing. I don't think I described it very well. But the idea being, having the salt is critically important for being able to perform a brute force. Without the salt, it's very likely no password on earth could create a collision because you just may not be able to - because the salt would be binary, and the password is ASCII. So it's a much smaller character set in terms of the bytes. ASCII is a much smaller range, very likely that you could never create a collision. So I would say, I mean, it's not crucial. But if you could put the salt somewhere else, that would be a good thing.

Leo: We have two more. Are you ready?

Steve: Yeah.

Leo: Question 10, Keith Rollin in Sunnyvale, California wondering how is overwritten disk data recovered? This is another one of those tinfoil hat things: In order to securely wipe a hard drive, operating systems and third-party utilities offer facilities for overwriting data multiple times. I've seen options for overwriting three, seven, 11, even 35 times. The implication here is that someone can still recover previously written data from a disk even if it's only been overwritten a few times. My question is, how is this done? If there are multiple images of a block in a particular physical location on the disk, how does the standard driver know which bits are the current ones? How does a recovery program know which bits are from one generation back as opposed to two generations back? How prevalent is the technology for reading old disk contents, and how many times should we overwrite a block to obliterate old data?

Steve: And I agree with you, Leo. I would now, more and more, I would classify this as an urban legend. That is, the idea that it is possible to go back multiple generations. This 35 times is the result of one famous paper on the Internet where someone analyzed very old generation drives, meaning MFM or RLL drives, where we knew how the data we were putting in was turned into flux reversals, because it's the reversals of flux on the drive which stores the data. With contemporary drives, it's amazing how many different stages, they go through at least three stages of data manipulation before, that is, the user data is transformed in three very different ways before the final flux reversal pattern is generated. Meaning that the user has diminishingly little control, based on the data they write, over what's finally written on the drive.

An earlier version of SpinRite, several generations ago, actually it's SpinRite 3.1 and 4, I think, did something where I reverse-engineered what were called the endecs, the encoder decoders of all the popular hard drives. I found them all, I reverse-engineered them, I knew how they related data to flux reversals. And I came up with this - it was called the "flux synthesizer," actually, which synthesized patterns based on the drive's make and model to do the best job of finding defects. And that sucker really worked. The problem is, you just can't do that anymore. In the same way that you cannot deliberately write data to create flux patterns on the disk, that's just - it's gone from all of our contemporary drives. Similarly, there's no way to do, like, worst-case data patterns. They existed once. They don't exist anymore. And error correction is so prevalent that it's not clear that that really matters, either, because error correction is always there, always in use, and it's solving problems like allowing us to skip over little defects and correct for them. My feeling is that two passes of pseudorandom data such that there's no record of what was written is all you absolutely would ever need. I would argue that one pass could be reverse-engineered, but not two. And even then it would take, I think, as you said at the top of this, Leo, it's another tinfoil hat issue.

Leo: It all came from a guy named Peter Gutmann who gave a speech in 1996.

Steve: That was Peter.

Leo: Yeah, asserting that this was possible, never demonstrating, never explaining, never showing any real evidence. And even then, what work they did required a scanning tunneling microscope per bit. It's just not - it's not - it doesn't - no.

Steve: No.

Leo: And it takes a lot of time. But it's become - there's so many things that are become revealed truths in technology because we don't understand it very well, and so we just all kind of accept it. And that's one of the revealed truths, that you must overwrite 35 times, or somebody somewhere could figure it out.

Steve: Yeah. The next major revision of SpinRite, which is a ways off because I'm going to do a minor one first, but it will incorporate the what I call "beyond recall" technology, and it will have an extremely good pseudorandom number generator, because I know how to write those now, having done all this crypto work, and it will very quickly do a secure wipe of the drive, and also do what it can about the safe area, the protected area, and the relocated sectors, a ways of getting into there, too. So that'll happen.

Leo: Our last question, ladies and gentlemen, is the Revelation of the Millennium. I need a gong or something. Ed in Cleveland says: I've solved all of the world's problems without even trying. Steve, I believe I've come up with a way to solve all the world's problems. I'm contacting you because I know you will be able to confirm if my theory is valid.

I will start with the simplest example: The recipe for tacos, when stored as a text file on a computer, is just a string of digits. If you randomly generated digits, you would

eventually produce a set of digits that is identical. You'd also produce a lot of noise, along with many variations of the taco recipe. If you never stopped, you would produce every recipe in every language that has ever existed. In fact, I'll point this out, you would produce everything ever, including all of Shakespeare. But let's go on.

You could eventually do the same with music. Start producing a random set of digits, and you eventually end up with "Hey Jude" by the Beatles. The process will produce the entire Beatles catalog, along with the most beautiful Beatles songs that John and Paul ever dreamed of making. These random sets of digits will even contain a recording of "Yellow Submarine" with the fifth Beatle, Steve Gibson, on lead vocals.

Steve: Now, there's a scary thought.

Leo: Ha ha. It's true, though.

Steve: Yeah, it is.

Leo: The problem is the word "eventually." But anyway, we'll get on. Among the text documents, sound recordings, and videos that could be produced by randomly generating digits is the cure for cancer, the end of war, the source of unlimited free energy, and everything you can and cannot imagine. I think Douglas Adams wrote about the Infinite Improbability Drive.

Steve: Ah, love that, yes.

Leo: This process could even produce a documentary video of future events; a video of the highlights of my life including the last breath I take, absolutely accurately and perfect in every detail. My point is, although it may not be practical, it is possible to produce a digital version of everything that has ever happened in the past and will happen in the future, and also to produce things that would never have been invented by man. Do you agree that this is possible?

I point you to, Steve, a very famous and infinite number of monkeys typing on an infinite number of typewriters, which was conceived of long before digital technology.

Steve: Ah, yes, the lure of digits, Leo.

Leo: It's no different 'cause it's digital.

Steve: Yes. I just got a kick out of this posting. And what it tells us, as listeners of this podcast who are interested in things like bit lengths and in crypto and in probabilities, all of which we work to understand and deal with, is that Ed, of course, is correct. And here we are looking at the impossibility of finding the 128-bit symmetric key used to encrypt

communications because, for example, that's what protects SSL. Yes, there's all of the fancy public key, private key, all of that. But that just negotiates, remember, the symmetric key, which is much shorter. And we can't try all of those combinations. There are too many in only 128 bits. And if we look at how few characters that is, then we get some idea of the number of bits required to represent one of the Beatles songs. And that, sure, you could change a few of the bits here and there, and the song wouldn't be drastically different.

But when you start doing the math and looking at how many combinations of bits there are, it's true that everything that has ever existed is in some pattern of bits. The problem is that so is everything that will never exist, and there's a lot more of that than there is - there's a lot more of the junk than there is of the good stuff.

Leo: The real problem is this term "infinity," which somehow implies that it's a number, that it is in the same category as one, two, three, four, and five. That if you kept counting long enough you'd get to infinity.

Steve: Well, and it's because we deal, we humans...

Leo: We don't get it.

Steve: Yeah, exactly. We operate in a world of how far can we throw this stone, and when will my car run out of gas, and physical, conceivable, real-world sizes and numbers. And what I like about this...

Leo: Infinity ain't one of them, by the way.

Steve: Right, exactly. What I liked about this is that it reminds us that combinations of bits are incredibly powerful.

Leo: Right.

Steve: I mean, he's right; Ed is right. All of those things exist as a combination of bits. The other thing this says to me is a little bit about entropy, which is we are highly ordered in our combinations of bits. Computers and the technology we're using, the fact that you and I are communicating in real time, 500 miles from each other, and audio is going back and forth, I mean, this is this incredibly supreme accomplishment against entropy, where think of all the things that could go wrong, all of the ways that these bits could be different, and nothing would work. Yet here it is, and we're talking to each other. It's fascinating. But, yes, the world is big; the universe is big. And as you said, Leo, infinity is a very large number.

Leo: It's a bigger number than one really thinks of.

Steve: It's a lot further than you can throw that stone.

Leo: And then, for further research, Zeno's Paradox, infinitely small numbers. And there's some who say in string theory that perhaps that thing that he was talking about is actually occurring, and that we are existing on one of the many strings of infinite possibilities, and we just happen to occupy one of them. I hope you're enjoying yourself.

Steve Gibson is at GRC.com. That's where you can find 16Kb versions of this show, which are mostly but not entirely random. And also, of course, you can get the transcript. You can get lots of free stuff. And the most important thing, SpinRite, the world's finest hard drive maintenance and recovery utility. You gots to have it. If you've got a hard drive, you need SpinRite. Steve will be back, not Wednesday, but Tuesday. I should tell everybody.

Steve: Yes.

Leo: We talked about this. March 7, next Wednesday, is Apple's iPad 3 announcement day. So we'll be doing MacBreak Weekly that day. We'll start our live coverage at 9:30 a.m. Pacific, 12:30 Eastern.

Steve: Be still my heart.

Leo: Yeah. So we're going to flip-flop you and put you in MacBreak Weekly's time, 11:00 a.m. Pacific, 2:00 p.m. Eastern.

Steve: Tuesday.

Leo: 1900 UTC on Tuesday, March 6. So...

Steve: And we will discuss SPDY, S-P-D-Y, finally, what Google has been doing to propose improvements to the most-used protocol on the Internet, which is HTTP, which all of our web browsers use - how to make it go more speedy.

Leo: Awesome. Steverino, thank you so much.

Steve: Thanks, Leo.

Leo: Have a great week. We'll see you next time on Security Now!.

Steve: Bye bye.