Listener Feedback #117

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-300.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-300-lq.mp3

Leo Laporte: This is Security Now! with Steve Gibson, Episode 300, recorded May 11, 2011: Your questions, Steve's answers, #117.

It's time for Security Now!, the show that covers your security, privacy online. And the man of the hour, as always, Mr. Steve Gibson of GRC.com.

Steve Gibson: Hey, Leo.

Leo: Good to see you. And we've fixed your camera.

Steve: Great to be with you as always. Episode 300.

Leo: Yeah, we should do something, party or something.

Steve: Yeah, well, that's not one of my favorite numbers, though. I mean, it's not 256 or 512. When we get to 512, baby...

Leo: You're going to like it, you know, we were trying to figure out what to - we're going to sell bricks in the new studio. People can put their names or their company name or their Twitter handle on it or whatever. And we decided to make the prices of the new bricks $128. And if you want the paver, the square paver, that's $512. If you want a copy of your brick, it'll be $64. So you'll like that.

Steve: Very nice. Those are good numbers, absolutely.

Leo: Yay. Let us - we have - it is nominally a Q&A episode, but you have thousands of things to talk about.

Steve: Well, yeah. So much happened this week, and some interesting things that I wanted us to take our time and talk about, that I was afraid if we tried to, like, rush through those in order to get to a regular large set of questions, that we would not do the front of the show justice, which I really, in this case, there's really some interesting things for you and me to talk about. So I had a couple of questions that were actually inspired by questions that I saw through Twitter, but which I've also seen in the regular Security Now! /feedback page. So I'm going to sort of wrap up some things at the end from previous shows, sort of in the Q&A format. But otherwise let's just talk about what happened this week because there was all kinds of stuff.

Leo: Oh, I love it. I love it when there's good news.

Steve: And I'm sure we've got a good podcast for everybody.

Leo: We will indeed. All right, Steve. Let's get down to work. We've got business to do. First of all, did you read, did you finish Russinovich's book?

Steve: Yes, a couple days later. You'll remember that I, like, made myself go to sleep Tuesday night because I was at 50 percent of the book and really enjoying it. And by a couple days later I had finished it. And so I'm glad you reminded me of that because I want to tell everyone that I recommend it without hesitation. As I was, you know, I was thinking about the surprise you expressed at, like, Mark can write. And if it said "Written by Michael Crichton" on the front, I wouldn't have thought twice. I wouldn't have thought he was turning senile or he'd lost his ability or anything. I mean, Mark can write, believe it or not, not only computer code and deeply understanding the technology of Windows. He's regarded, as we know, as probably the person who's more about Windows security than anybody else there is. I mean, more than Microsoft employees because he was on the outside poking around and creating all those cool utilities back in his Sysinternals days, before Microsoft bought them.

Anyway, the book is, as you'd expect, is factually accurate. And I wish I knew that everyone who would read it had read it, so I wouldn't be concerned about spoiling it. But the title is "Zero-Day." And so from what I said last week, we can imagine that this is Mark's very interesting, really gripping portrayal of how it could be that something really bad happened on a global scale. So it's, I mean, it's really interesting. And as I said last week, it felt to me like our podcast, a fictionalized version of the stuff we talk about. So my sense is our listeners in particular, probably more than any other audience I can imagine, would just, if you're a book reader, if you like fiction, and you listen to this podcast, I mean, there's the criteria. I recommend it without hesitation. It was just very pleasant and not Hamilton length.

Leo: I'm in the middle of "The Dreaming Void," and I know what you mean when

you say "Hamilton length." Oh, my god. But, you know, Dr. Mom read this book, and on your recommendation. She's a regular in our chatroom. And she said the same thing. She said, "I missed, like, a good night's sleep. I couldn't...." And, by the way, she said she was terrified by it. And that's, I think, interesting.

Steve: Actually, Leo, I found myself thinking, I don't want the bad guys to know this stuff.

Leo: Yeah.

Steve: I mean, he went further than we had. One of the chilling things he does is he reminds us of how pervasive our use of computers has become in this book...

Leo: We rely on them entirely.

Steve: ...by citing specific examples. And, for example, there was, in one case, and this isn't giving much away, but I guess we have an advanced plane, the 787, which not only does the computer completely fly it, so that there are no more cables, I mean, it's all fly by wire, but computers designed it. So the computers are designing it and then flying what it is that they designed. Well, what happens if something goes wrong with that? So, I mean...

Leo: Come on, Steve.

Steve: And on and on and on. So, I mean, I know what Liz meant. It is, as I'm reading this, I'm thinking, well, this just tells them how to do this. I mean, he goes further than we have gone in the show. And it's wonderful, so.

Leo: There's actually, in that light, the fly-by-wire airplanes, there's actually a debate. There's one airline, which shall remain nameless, that its pilots' union insisted that a pilot be able to do anything, completely override that airplane. And the other - what happens as a result is pilots can do things that will crash a plane. And so the real question is, you know, I guess you've got, you put a human there for a reason. They should be able to override. But should a human be able to do anything is the question.

Steve: Well, and in fact it sounds to me like Mark did a lot of research for this because, I mean, he talks about exactly what you're saying. He talks about the psychology of the pilots and how in demonstration after demonstration during their training, they were shown that no matter what situation the plane was in, it would do a better job of recovering than they could.

Leo: Doesn't that sound like "War Games" or "Titanic" or, I mean, haven't we seen

this movie? We've seen this movie.

Steve: Yeah, yeah. Well, and, I mean, so...

Leo: He has a trailer, by the way, on his website now for this book.

Steve: I'm glad.

Leo: I can't wait to read it.

Steve: I just - I would love to be able to say more. But it would - it's not fair to the people who like the idea of, like, the kind of things we talk about in this podcast, fictionalized by, like, the leading Windows security person in the world, who also can tell a story and develops characters. I mean, now sitting here, I know all these people that Mark created and their interrelationships, I mean, vividly. I mean, he's really a fiction writer.

Leo: Well, but it's, you know, they say write what you know. He sure knows his stuff.

Steve: Yeah, but it wasn't just...

Leo: Oh, I can't wait to read it.

Steve: It wasn't just dry, like a recitation of security nonsense. I mean, there's a full plot. There's, you know, I just can't say anymore. But, I mean, there's...

Leo: I'm going to get it. I can't wait.

Steve: There's bad guys. There's murder. There's all kinds of stuff. I mean, it's really - there's assassins and, oh, it's just wonderful. So, yeah. I'm so pleased to have read it and to be able to tell our listeners because I'll bet you, I mean, I know that our sci-fi recommendations - because you have followed me with Hamilton and everything, and of course Michael McCollum at SciFi-AZ. I get so much feedback from people, it's like, oh, I'm so glad I got turned on to this. So here's another one. I don't know how much more Mark will crank out. But I'm reading No. 2 if there is one.

Leo: Zerodaythebook.com is the website. You can buy it there. And it's always a good idea to buy the book through the author's website because he'll get, in addition to royalties, usually he'll get like an Amazon...

Steve: Nice, what do they call it, affiliate, right.

Leo: ...affiliate fee. So it could mean a buck fifty or more, per book more for him.

Steve: Yes.

Leo: Kindle, did you read it on the Kindle?

Steve: I did, yes, because I'm just a Kindle fanatic. And I did get a tweet from Steve Wooding in Hampshire, U.K., a listener, who was disappointed to see that it wasn't available outside the U.S. on the Kindle store through Amazon. So maybe it just hasn't happened yet. I can't imagine why...

Leo: Publishing is such a byzantine and medieval system.

Steve: You think it's the publisher who says...

Leo: Oh, it's always the publisher. It's publishing rights. It's just - it always is. There's no audio book probably for the same reason. And sometimes the audio books are available in one country, not the other; Kindle editions in one country, not the other. It's because the U.S. publishing rights are different from the overseas publishing rights.

Steve: Well, and we do know, for example, that some Kindles will allow themselves to be read in their audio mode, and it blanks, is blocked in many others.

Leo: Great reviews, too, by the way.

Steve: Oh, really?

Leo: Fifty-two five-star reviews, yeah.

Steve: Oh, I'm really glad. Like I said, if there's this, like - it's exactly like your reaction. There's nothing, as I'm reading it, and even as I'm reading the second half after hearing your comment last week, your surprise. I mean, and I was also...

Leo: Well, you don't expect a guy like this to be a good writer, to be honest.

Steve: I know. I mean, I'm sorry, but...

Leo: And look at Michael Crichton, who was a physician first. I mean, it's not unusual. Scott Turow, who is a lawyer first. You don't have to be born and bred to be a writer to be a great writer.

Steve: Well, and Crichton, a lot of Crichton's work was writing what he knew.

Leo: Right.

Steve: Because he was an M.D., and he gave us "Andromeda Strain," and he understood the biology and the molecular level thing. Well, here Mark has done the same thing. He's given us a book which would satisfy any listener on this show, I mean, in terms of what he tells us. And, yeah, I can understand Liz being a little bit frightened. It's like, it is, I'm thinking, whoa, it's too possible, the scenario he paints.

Leo: The chatroom is saying Asimov was a biochemist. John Grisham was also a lawyer. So a lot of lawyers become writers because nobody wants to be a lawyer. Seriously. The happiest lawyers I know are all ex-lawyers.

Steve: Well, and Grisham's books also...

Leo: He knows the law.

Steve: Exactly. Exactly.

Leo: Mary Jo Foley gave it five stars on Amazon, as well. So, yeah, okay, sold.

Steve: Yeah.

Leo: That's great.

Steve: Oh, they're Amazon reviews. Oh, good, I'm going to go back. I thought - I didn't know where they were.

Leo: Give it an Amazon review, yeah.

Steve: Okay, good. I'll do that when I'm through with the podcast. We had a blessedly small Patch Tuesday just a couple days ago. There were only two updates. One contained two fixes; the other contained just one. So, which is nice. We get to take a breather following last month's once again record-breaking Patch Tuesday. And one of them was critical. But really nothing to worry about for typical end users. It only affected enterprises who are the typical users of Windows Server 2003 or 2008.

And it was a vulnerability that was found in the Windows Internet Name Service, WINS, which is not even installed by default. This was an early protocol that Microsoft used back in the dawn of the Internet, and DNS essentially replaced it, the Internet's name service, domain name service. So it's not installed on machines. It's a problem if you did have an enterprise system using Server 2003, 2008, and had the WINS service exposed to the public, that is, to the public Internet. That would be a problem.

So this was disclosed responsibly. Microsoft patched it. And they rated it critical because, if you fit the scenario, then you're in big trouble. You need to get thing fixed. But most people aren't going to be. And then the other two fixes in the other one patch were two things in PowerPoint. And I don't know why Microsoft rated them important instead of critical because they were remote code execution. You go to a website that launches a PowerPoint on you, and it can run code. But for whatever reason Microsoft - oh, I guess it's that you did get prompted. There was a - it wasn't an automatic thing. You would get prompted by PowerPoint, do you really want to run this? And there wasn't a way around that. So I think that's what allowed them to drop it down a little bit. So anyway, worth doing, as always. But way more tame than the patch blast that we received the month before.

Leo: Yeah, that was a record.

Steve: Yes, it was. They broke, well, in order to keep breaking your record, you have to keep doing worse and worse.

Leo: Yeah, we don't want a record. We're not gunning for a record.

Steve: Let alone a succession of records where each one escalates the prior one's record. Now, this is a really interesting development, I thought. TOR, which we've talked about, The Onion Router - and I know that we're picking up new listeners all the time. So if you don't know the acronym TOR, which is The Onion Router, it's definitely worth going back and listening to the podcast we did on it [SN-070]. The technology that was developed - and this was developed with the help of funding from our friends at the EFF, the Electronic Frontier Foundation, that are powerful privacy advocates. The technology is such that it's an anonymizing system, which is encrypted and anonymizing, so that when your traffic gets out on to the Internet ultimately, it's impossible, I mean really, truly impossible, and I don't say that often, to know where it came from.

And the way it's done is that there's a series of hops - I'm going to just summarize the podcast, but really go back and listen to it if you want more - a series of hops that your traffic makes among TOR servers, essentially, or like TOR routers, that are - there's a large network of them. So your particular traffic will choose some number of hops and routes, and you have control over that at your end. Each of these routers publishes a private key, I'm sorry, publishes a public key, has a private key, publishes its public key.

Knowing then the order in which your packet traffic is going to hop among these routers, you successively encrypt your data, one after the other, in the reverse order that your traffic is going to be hopping among the routers. So you first encrypt with the last router's public key. Then you encrypt with the next to the last router's public key. Then you encrypt that. So you've got multiple levels of encryption, thus this concept of an onion, you know, the notion of layers of an onion.

Leo: You peel it back, yeah. But you also see there the negative of this, which is there's considerable overhead.

Steve: Yes, yes.

Leo: Slows you down a little bit.

Steve: So you then send this big blob, this onion, off to the first router. What's cool is that it has its private key. Only it can decrypt what's inside that layer. But it can't decrypt anything more because that next layer is encrypted with the next router's public key, can only be decrypted with its private key. So all that router can do is take a wrapper, a shell off the onion and forward it to the next one. And it can do the same thing. It can only take off the outer wrapper that it receives and then forwards it again. So at every stage the intermediates can never see your traffic until you get to the final endpoint, where that final router takes the encryption that it knows how to remove, which was, because it's the last router, that was the first encryption, finally decrypting your traffic, and it puts it out on the Internet.

So the point is that this is a big network of these. All of the traffic among them is encrypted. Oh, and I should mention also, the next router hop and all the IP stuff is also encrypted. So it's not even the case that an earlier router knows where it's going to go after it goes to the router it sends it on to. That's contained successively within these layers. So it was really designed beautifully to allow people the freedom of accessing the Internet anonymously. The problem is that you have, to actually use it, you need to download the TOR system and a client. Then, for example, if you're using Mozilla, Firefox, you need to use its add-on in order to link to that. And there's a bunch of configuration stuff. And the point is it's all a lot of overhead.

What's happened is the TOR folks have recently just announced that they're going to fork the Firefox project from Mozilla and integrate all of the TOR technology into their own Firefox.

Leo: Oh, my goodness. That's interesting.

Steve: Well...

Leo: So you'll have a - you'll essentially have a browser that does it automatically.

Steve: Yes. Yes.

Leo: Wow.

Steve: And they're excited because it dramatically simplifies the use of TOR. I mean, I've never bothered because I don't really have anything. But it's like, I could see having a copy of that around, if for some reason - like I'm doing some security research, and I

wanted to poke at something but not leave my identity there. The other thing, because it would be essentially a TOR/Firefox browser, is they can do other anonymizing things to the headers. And, I mean, you can imagine it would turn on the DNT. I'm just joking because you couldn't be tracked through this anyway. But...

Leo: Wouldn't that be funny.

Steve: It would give them - I guess what's happened is there have been times when the TOR folks have been a little frustrated that Mozilla's priorities have been different from theirs. Or things weren't fixed as quickly. Or they've been waiting for Mozilla to do things that they wanted. And so they just decided, okay, we're going to fork the project and integrate TOR into this browser, which...

Leo: That's the beauty of open source, you know? That's why we love open source.

Steve: Yeah. And it affords this. And so, anyway, I'll keep an eye on this. It's been announced, but it doesn't exist yet. So as soon as it does exist, I will be sure to let our listeners know because it would just be cool to have this. It'll be a lot slower, as you said. There's a lot of overhead in doing this. But just to have a browser as simple as opening Firefox, a state-of-the-art browser, which automatically uses TOR and deals with all of the overhead, knows how to access the network, and where you know that you are absolutely untrackable.

I guess what I like it about it, too, is that there are certainly so many people who have a bona fide, useful, sort of innocent interest in doing this that can't handle the technology. All of the TOR downloading and configuring and add-on and all that, that's enough off-putting that they don't end up getting the benefit of it. So now it'll just be much easier.

Leo: By the way, we have a whole show on TOR [SN-070]. If you want to know more about TOR, go back to the archives because there's a lot to talk about with TOR. And there is one, there's a little bit of a vulnerability in the sense that the endpoint can be compromised.

Steve: Yes, I'm glad you said that because it's worth mentioning here in this discussion, although we did talk about it then. And we've talked about TOR from time to time since. And that is that the so-called "exit nodes" of the TOR network, you can imagine that law enforcement would be very interested in the traffic to and from those specific points. So in the same way that any VPN provider, you wonder, if you're routing all your traffic through a VPN provider, well, they end up being someone that bad guys, or even good guys, law enforcement types, would be interested in making - in, like, surveilling because there's some reason people are using TOR. There's some reason people are using a VPN. Hopefully just for security, but it could be for nefarious purposes, as well. So you're right, Leo, there is that side to it.

Okay, this is weird. We know that we've got problems right now with uprising in Syria. Apparently the Syrian Telecom Ministry has been perpetrating a nationwide man-in-the-middle attack against their citizenry for Facebook. What they're doing is - the EFF got involved. Someone named Mohammad, who is a Syrian, inside Syria, brought this to their attention and shared his logs and was able to dump his certificate out in order for

them, for the EFF to take a look at it. And what they found was that what Mohammad had was the real IP of Facebook. Which means that this was not using DNS spoofing or tampering in order to give him the wrong IP. He had the right IP, which meant it had to be using routers and proxies, which means that someone is actively filtering the traffic in and out of Syria. And in the case of HTTPS access to Facebook - and by the way, we'll talk a little bit about that later - which Facebook reports that 9.6 million users are now using Facebook over HTTPS, thanks to the fact that they've added that option that we've talked about, the persistent use of SSL encryption over connections to Facebook. So somebody is intercepting that traffic and returning an unsigned Facebook certificate.

Leo: Oh, wow.

Steve: So you get warnings. I mean, any browser that sees a certificate claiming to be from Facebook which has not been signed by a certificate authority that the browser trusts will pop up a security warning.

Leo: Of course.

Steve: But many users don't really understand what that means, and so they click past it. And in clicking past it they're allowing essentially their government, in this case the Syrian government, to monitor - to decrypt and then reencrypt in a proxy, in a transparent proxy, all of their traffic that passes through in order to monitor what they are seeing and posting on their supposedly secure Facebook page. Obviously there's no security here. So it is a concern that many governments do control certificate authorities. I don't know, apparently Syria doesn't and couldn't talk anyone into giving them a certificate for Facebook. Which is a good thing because that would be bad to have...

Leo: Because then you wouldn't know; right? It would appear to be a...

Steve: Well, actually, anyone who was using the very cool certificate auditing tool I talked about last week, I'm trying to remember the name of it, just a minute here, let me look, add-on, it's at the top of my list, I remember. Certificate Patrol. I'm still using Certificate Patrol and really like it. So anyone using that would be notified, even if Syria had a properly signed certificate, because Certificate Patrol, before Syria did this, and this has only been done recently, if you had ever connected with Facebook prior to that time, then Certificate Patrol would have cached the certificate that you had. Then when Syria tried to do this, well, first of all, not only - oh, yeah. If Syria tried to do this with a valid signed certificate, Certificate Patrol would say, wait a minute, this doesn't match the certificate you had last time. And it would present them both to you and allow you to then look at what's going on. And, I mean, it does some nice forensic analysis.

So it would say "This certificate is not signed by the same certificate authority that signed the Facebook certificate you were using before." It's unexpected that the site you are connecting with would change certificate authorities. I mean, certainly it can be done. Actually I'm planning to do it because I think I'm going to drop VeriSign after seeing that DigiCert is signing so many, well, in fact they're signing Facebook's.

Leo: Really. That's the one that's less expensive.

Steve: Yes, way less expensive...

Leo: Awesome.

Steve: Yeah, I mean, so I could afford...

Leo: If it's good enough for Facebook...

Steve: I could do EV certificates. Yeah, if it's good enough for Facebook...

Leo: How much are they?

Steve: I think they're down to, like, $300.

Leo: Oh, we really should do it for TWiT, then.

Steve: Yeah.

Leo: We are going to start, with the new website, we will allow people to log in and post stuff and so forth. So I think we should absolutely do that, yeah.

Steve: Yeah. And for me, too, GRC, I can't, I mean, VeriSign is thousands. And it's like, oh. And it's not just thousands once. It's thousands every time you have to renew, which we know is a good thing because it does keep bad certificates from never expiring, and they need to. But, boy, they're just raking in the dough, so.

I wanted to bring to our users' attention an interesting study that some researchers in Belgium and France did. This was actually picked up by the Register over in the U.K. But there are many sites, I guess like a hundred file hosting sites - like RapidShare, FileFactory, Easyshare - which allow users to upload files of any size, typically large files, and then they give those files a unique URL, the idea being that, I mean, the URL almost looks sort of like a little crypto key. It's not short like a little bit.ly key where the whole goal is to have it short. It normally looks like it's very unguessable, and that's the point is they're saying that you don't have to worry about anybody else getting a hold of these files because, look at that URL, no one's ever going to figure that out. Well, these researchers conducted a couple of experiments. They put some of their own web spiders on these sites and collected thousands of private files. And they then...

Leo: Wow.

Steve: Yes. So...

Leo: Just by randomly guessing...

Steve: Well, or maybe they looked at, like, at a succession...

Leo: Maybe they found an algorithm.

Steve: Exactly, they looked at a succession of the URLs and saw what the pattern was. And it's obviously not very good because, if it were really good crypto, and it was sufficiently long, then the chance of just by brute forcing, basically they're brute forcing the URLs. The chance of brute forcing a sufficiently long URL would never make it worthwhile. You just...

Leo: There are many of these services. Do they say which ones?

Steve: Well, they said a significant percentage of the hundred file-hosting services they studied made it trivial for outsiders to access the files simply by guessing the URLs that are bound to each uploaded file. And RapidShare, FileFactory, and Easyshare were mentioned. And so the second thing these researchers did was they put their own files up with beacons in them that would allow them to determine if somebody else downloaded them.

Leo: Oh.

Steve: They never gave out the URLs. And quoting from this, "They also used the sites to store private files that contained Internet beacons, so they'd know if anyone opened them. Over a month's span, 80 unique IP addresses accessed the so-called "honey files" 275 times.

Leo: Oh, dear. So that's a reasonable thing is, well, one thing if we can brute it, but is anybody trying this? And apparently they are.

Steve: Yes. Which if course indicates "the weakness is already being exploited in the wild to harvest data many users believe is not available for general consumption." So I wanted to give a heads-up to all of our listeners that this is an instance where you pre-encrypt what you are going to post up there. And then through a separate channel you let your recipient know the passphrase for the encryption so that you could still use these file hosting services. But don't just put a spreadsheet of important financial data up there and assume that no one is going to ever get to it or see it. We have evidence now that these sites are being, literally, they are being dredged for information that people would not want to have made public.

Leo: That just tells you something, wow. Wow.

Steve: Uh-huh.

Leo: Isn't that interesting. I don't, you know, I use Dropbox and other stuff. But I would not, I agree, I think it would be imprudent to use this for something private. But I wouldn't always assume that, you know?

Steve: I have a new acronym. Unfortunately, it's PEE, P-E-E. I can't think of anything else.

Leo: That's fine.

Steve: It's Pre-Egression Encryption.

Leo: Oh, Pre-Egression Encryption.

Steve: Yeah. So you pre-encrypt before anything egresses from your location.

Leo: I think that's a good acronym. I certainly won't forget it.

Steve: Pre-Egression Encryption.

Leo: Wow.

Steve: So, okay. I loved this next story because I couldn't have - it sounds like someone's been listening to the podcast who did this report. First of all, it's based on a new bill that has been submitted by Sen. Rockefeller. And this is the Do Not Track Bill, legislation that's the thing I've been saying we need. It is the enforcement of the DNT header, essentially. And this was a story carried by Cecilia Kang in the Washington Post. I just want to read this because, like I said, I couldn't have written this any better:

"Sen. John D. Rockefeller ... on Monday introduced an online 'do not track' privacy bill that would allow consumers to block Internet companies from following their activity on the Web. The Do-Not-Track Online Act of 2011 comes amid increased attention by lawmakers on creating privacy rules for the Internet. The White House has called for such rules but has not supported a specific mandate that would block companies from tracking users.

"Rockefeller, chairman of the Commerce, Science, and Teleport" - teleportation. I wish.

Leo: I wish we had teleportation, yeah.

Steve: "...and Transportation Committee" - I've been reading way too much science fiction - "said in a statement that recent reports of privacy breaches show that companies have too much freedom to collect user data on the Internet. His legislation would force companies to abide by a consumer's choice to opt out of such data collection. The Federal Trade Commission would draw up specific 'do not track' rules. The agency and states' attorneys general would enforce the law. And the legislation would apply to mobile phones a growing platform for accessing the Internet."

It says, quote, "'I believe consumers have a right to decide whether their information can be collected and used online,' Rockefeller said in a statement. 'This bill offers a simple, straightforward way for people to stop companies from tracking their movements online.' Already, Microsoft's [forthcoming version of Internet Explorer] and Mozilla's Firefox browsers have been redesigned to allow users to block marketers from tracking what sites they visit and their other activities online." Now, we know that's not quite true because they don't actively block, they simply request. And thus has been the controversy of the DNT header, is because everyone says, yeah, but it's optional. People can ignore it.

And it says, so going on in the article, it says, "But without a law, no Internet company is required to honor the consumers' request, privacy groups said. 'This bill will put regulatory support behind these industry initiatives and make sure that online providers listen to the many consumers who want to clearly say "No" to online tracking. 'This complements the comprehensive online privacy legislation introduced by Senators Kerry and McCain last month.'"

So I just wanted to bring this to the attention of our listeners. This is, I mean, it is what we need in order for the DNT header to happen, essentially. And...

Leo: Now, it would not prohibit a website from tracking a incoming request's IP address, would it?

Steve: No. And Leo, I've been thinking a lot about what you said when we talked about this, and TWiT. And when you look at the detailed legislation, and I have, although I'm not going to bother our listeners with it because we know that at this point it's just been a bill that's been submitted. It's been talked about being set up, being hooked onto Kerry and McCain's bill as an amendment to it in order to incorporate this technology. But they all really do talk about information gathering. And, I mean, like address and age and specific criteria, rather than just like using their IP address in order to disambiguate queries. So I really think that what Podtrac is doing and which you and your industry depend upon for developing good numbers, I think that's going to be okay.

Leo: Well, and not just me, but every website in the log keeps track of an incoming IP address.

Steve: True, true.

Leo: I mean, that's just what happens. And I don't know how you would turn that off. That's how browsers work. Otherwise they can't have a conversation.

Steve: True. And I guess the problem is, when I was talking about, well, it'd be easy to just bring up a dialogue box to say we're Podtrac, we need to count you in order to credit TechTV with your download. The problem is iTunes is downloading podcasts autonomously. So it's unable to do that. Yeah, I just don't - I don't think it's going to be a problem. I mean...

Leo: I hope not, for me. I don't think either because I just think it would break the Internet if you said websites cannot record an IP address of an incoming request. That would just break websites. It would break everything.

Steve: Yeah. Well, and remember we're talking, there's certainly a difference between first party and third party. There's a general understanding, well, in fact Google and Facebook are screaming at the same time about newly introduced California legislation. This is a Senate Bill 761 that was introduced by Alan Lowenthal, who's out of Long Beach. He's got legislation which is proposing that companies doing business, doing online business in California would be required to offer an opt-out privacy mechanism for consumers. And, I mean, and frankly, there's a long list of people who signed a letter objecting to this - Google, Facebook, Yahoo!, American Express, Experian, Allstate, Time Warner Cable, the MPAA of all people, Chamber of Commerce and, like, 20 other companies have said whoa, whoa, whoa. Actually they said in their letter that "Senate Bill 761 would create an unnecessary, unenforceable, and unconstitutional regulatory burden on Internet commerce. The measure would negatively affect consumers who have come to expect rich content and free services through the Internet, and would make them more vulnerable to security threats." I don't know how you get some of that. I mean, that seems like it's - their response is being overblown. But clearly there's a difference between people going to a site and third-party tracking, third-party data aggregation.

Now, the problem is, of course, that Google's existence is owed to the fact that they're able to, well, especially with the purchase of DoubleClick, they're able to serve ads that are interesting to people. I think it's been clever that they're able to use the search terms the user is querying in order to choose ads. There you're not really having to build a history or a profile, as many of these other media companies are. So again, and I think that what Podtrac is doing is just totally benign compared to this. But, well, and the good news is we'll end up with something somewhere in between that the legislators are happy with, and that the Googles and Facebooks of the world are saying, okay, we can live with that, too. I don't know what the answer is going to be. But there really is, really is growing tension on the Internet over this issue. It's not going away.

Leo: Well, I'm tense. I would, you know, okay. I don't want to go out of business. We get a lot of free stuff based on advertising, including all of the things you listen to on TWiT. And I don't want to go to a paid model. But advertising does require - it's a quid pro quo. I think...

Steve: It needs to have numbers.

Leo: Yeah. And I think that - I don't think - I think people understand that. I would hope. But the problem is, if it's opt-in, that's problematic.

Steve: Yeah. Well, okay. Now, in your case, you're not interested at all in who your listeners are. You're interested in not double counting them.

Leo: Right.

Steve: But you don't care what their demographics are...

Leo: We don't crack them.

Steve: ...their age range, how many children they have in the household. That's the kind of stuff which really creeps people out. When knowledgeable people look at the extent of what this data aggregation ends up meaning, and the fact that it ends up being de-anonymized ultimately - remember we went through that period of time where you'd, like, you'd have free offers, and you'd sign up for something? Well, those free offer sites were using advertisers and providing all of that non-anonymous information behind the scenes back to the advertisers and being paid heavily for it because it meant so much. So, I mean, that's what creeps people out. And I think that's what we're, I mean, that's the tracking. It's not just, hey, I went to a website, I want it to forget that I was there. It's, no, it's like, I want DoubleClick not to know where I've been and all the sites I've been today.

Leo: Right. Interesting. We live in interesting times.

Steve: We really do.

Leo: You know, it's always a risk when you start a new business, especially in a new area like this. I just read an article about a California state law that essentially puts money transfer companies out of business, if you're not doing it with a credit card, because you have to get a license from each state. And each state can be lots of money. Half a million dollars in the state of California.

Steve: Well, and iTunes blocked Bitcoin because they said...

Leo: Yeah. That's going to put Bitcoin out of business.

Steve: They said they were an inter- that Bitcoin was an intermediate currency, and that ran afoul of some iTunes regulation. And so there will be no Bitcoin app for the iPad or iPhone. Which is annoying.

A hacker, I guess he has to be a hacker because of what I'm going to explain here,

named Gordon Maddern, who has a site called PureHacking.com, he wrote, and this ended up being a big story: "About a month ago I was chatting on Skype to a college about a payload for one of our clients. Completely by accident, my payload executed in my colleague's Skype client."

Leo: Oh, boy.

Steve: "I decided to investigate a little further and found that the Windows and Linux clients were not vulnerable. It was only the Mac Skype client that seemed to be affected. So I decided to test another Mac and sent the payload to my girlfriend. She wasn't too happy with me."

Leo: I pwned her computer.

Steve: Although if she's his girlfriend she already knows what he's up to. "She wasn't too happy with me as it also left her Skype unusable for several days."

Leo: Wow, wow.

Steve: "At this point I figured out what was needed to execute code. So I put together a proof-of-concept using Metasploit and Meterpreter as a payload." Meterpreter is something we've never talked about before. It's a piece of Metasploit that allows for the easier creation of DLLs for injection into compromised processes. So it's like, okay, fine. So, he says "Lo and behold, I was able to remotely gain a shell" on the remote Mac. "So after a lot of trouble trying to find the right person in Skype to notify" - I don't know whether that's going to get better or worse, Leo. We haven't mentioned on the podcast yet something that I'm sure everyone probably already knows.

Leo: Call Steve Ballmer. He could fix it.

Steve: Yeah, that Microsoft bought Skype for $8.5 million.

Leo: Billion, yeah.

Steve: Yeah, billion. I will say that I'm happy that we'll be moving to Vidyo.

Leo: Yeah, I mean, I don't know if Microsoft will be a bad steward or anything, in fact. And we'll still use Skype when we can't use Vidyo. But it's nice to have alternatives.

Steve: Yeah.

Leo: And I think Skype will be around.

Steve: Although you've got to wonder if they'll be as diligent about worrying and managing and maintaining the other platforms.

Leo: One would hope.

Steve: Yeah. Anyway, so he says, "After a lot of trouble trying to find the right person in Skype to notify, I was able to get the correct details for the security team in Skype. I notified them on the security vulnerability and was given the standard, 'Thank you for showing an interest in Skype security. We are aware of this issue and will be addressing it in the next hotfix.'" And he says, "That was over a month ago, and there has still not been a hotfix released. The long and short of it is that an attacker needs only to send a victim a message, and they can gain remote control of the victim's Mac. This is extremely wormable and dangerous. Pure Hacking...."

Leo: Yow. I don't like the word "wormable." I don't know what that means, but it's not good.

Steve: And that was it. That was what the news covered was Skype for Mac is wormable.

Leo: Wormable.

Steve: What it means is that a non - because many people leave their Skype clients running all the time, and Skype up, and when you are in Skype you can see all the other contacts, this could be a flash worm that would run through Skype and in a matter of minutes take over all the Macs that are interlinked through Skype, is what we're saying.

Leo: That's really amazing. They didn't force an update, but they did offer an update.

Steve: Well, he says, "Pure Hacking won't give specifics on how to perform this attack until a patch from Skype is released. However, we will give a full disclosure after Skype takes action or a reasonable, responsible disclosure period has elapsed." Now, in his own update he has confirmed that Skype has fixed this issue in 5.1.0.922. But you have to go, I guess...

Leo: And push it. No, you have to actually say I want an update.

Steve: Well, no, I tried that. And it didn't give it to me. I'm back on .914. And I fired up my Skype, and I said check for updates, and it said no updates are available.

Leo: Huh.

Steve: So I think you have to go to Skype and download new Skype in order to get it. It didn't - not only did it not offer it, but when I said check for updates, it said nope, got none. So...

Leo: Well, I'm at 922, so I can't test it over here.

Steve: Oh, no kidding, so you're already updated.

Leo: Well, I did it, yeah. Soon as I read that story I updated it. Are you kidding? I run Skype. Now, this is only Skype for the Mac, by the way.

Steve: Wait, you don't want to be part of the flash worm that...

Leo: No, no, thank you. But this is, again, only on the Mac. And for those who are curious, we don't run Skype - we run Skype on Windows. For right now that's what you're on, for instance, is a Windows instance. But many of our users, many of our hosts use Skype for the Mac. And most of them have actually downgraded because they hated 5 so much that they've gone back to 2.8. So, which is, by the way, safe.

Steve: They did mess up the UI on...

Leo: Oh, did they screw it up.

Steve: Yeah.

Leo: That's interesting. And now, on my other Mac, where I'm running 2.8, it says "New version of Skype available," and it's going to 935. So...

Steve: Wow, okay.

Leo: ...who knows now what's happening.

Steve: So it sounds like they probably did push out something quickly because they were aware that having a flash worm through interconnected Macs on Skype...

Leo: Not such a good thing.

Steve: That would not be good, no. No. In fact, it may have lowered their purchase price to Microsoft.

Leo: I don't know. I don't know how it could get any higher.

Steve: Oh, wait, yeah, it took out all the Macs. Huh, maybe that wouldn't be - well, anyway. So. Facebook applications...

Leo: What? Yes.

Steve: ...turn out to have been accidentally leaking access to third parties for all time.

Leo: Again.

Steve: Uh-huh. Okay. So this was revealed by Symantec, who provided an analysis. They said, "According to Symantec's analysis, the problem was caused by a flaw in the old Facebook API which apps use to authenticate their account access. When a user grants account access to a web app" - like a Facebook app - "the app is given an 'access token' which it is then able to renew. Symantec said that this access token can be mistakenly inserted into a URL returned by Facebook" and provided to the app server, which then receives it. And so it will - okay. So I got myself tangled up here. "If the app loads an ad banner or analytics code as the next step, it will send that URL, which includes the access token, in the referrer field of its HTTP request for the content. This referrer data is likely to have been stored in the log file on the advertising or analytics providers' server."

So we've talked about referrer headers before, and this is, like, this is a classic "oops" privacy leak. So what happened is a Facebook app would identify itself to Facebook and say I need access to this user's account on their behalf. And the user has given the app permission, which the user would have. But the user has been assuming that the app would not leak that permission. And so what happens is, if the app then receives the URL containing the access token, which it needs in order to impersonate the user, if that app then showed an ad, then as happens with referrer fields, the referred by essentially would be the URL that the app used which contained the token which allows the account to be accessed.

So essentially apps have been leaking an impersonation token that allows third parties, advertisers and analytics companies, I mean, apparently there are, like, logs out there with these HTTP referrer headers in them, because that's one of the things that web servers log. And these never die. These don't time out. They don't get stale. This is a cryptographic token that allows anyone who gets it to impersonate that Facebook user. So all Facebook users have to change their password. That's what this comes down to.

Leo: What?

Steve: Yes.

Leo: Like now?

Steve: Like now.

Leo: What if I don't use apps? I don't use any apps.

Steve: Then you're okay. Then you would not have had...

Leo: Of course, how do you know what's an app?

Steve: Exactly. Any of these little goodies. Apparently it's, if something says they want permission to act on your behalf, they want access to your wall, they want access to your whatever...

Leo: Yeah. I use a lot of that.

Steve: Okay, well, those are apps. So...

Leo: But that token's a one-time use token, or is it a permanent token?

Steve: It's a permanent token.

Leo: Oh, fudgsicles.

Steve: I know. Now, okay. So...

Leo: That means these will all be logged out, too; right? All these apps I'll have to reauthenticate? Or no?

Steve: No. They will have to reauthenticate. But they're doing that anyway.

Leo: Oh, okay.

Steve: So what happens is, if you change your password, then you instantly obsolete all of this past leakage such that the app has to come back and get a new token and...

Leo: Why isn't this a huge news story?

Steve: I know.

Leo: Why isn't this on the front page of USA Today? They have 600 million users, all of whom were just compromised.

Steve: Yeah.

Leo: What can they do with this password? What can they do with this?

Steve: Anybody who has one of these tokens, and Symantec's blog makes it very clear that server logs all over the world are full of these tokens, are able to impersonate the Facebook user. They're able to do anything that the app, that the user gave the app permission to do. That permission has then been leaked.

Leo: I can't believe there's not a big banner on Facebook that says you have to change your passwords. I don't understand. I'm baffled by this. That means, I mean, I'm thinking of my son, every one of his friends in high school. They're not going to know about this.

Steve: No. Now, maybe the reason this hasn't been made more of is that, well, first of all, Facebook would rather not.

Leo: Well, yeah.

Steve: They have fixed the problem, and they've also moved to a new authentication system. They'll be using what's called OAuth 2.0, Open Auth 2.0. We're going to have to do a podcast about this in detail.

Leo: OAuth is incredible. I love OAuth.

Steve: Yes. Stina and I talk about it every time we get together for coffee, from YubiKey fame.

Leo: You're such nerds.

Steve: Ooh, let's talk about OAuth. So the good news is that Facebook is really tightening things up. They have implemented OAuth 2.0.

Leo: Good.

Steve: It's not mandatory yet, but they are going to sunset.

Leo: It would have to be mandatory with the apps, not with you, but with the apps that are requesting your authentication.

Steve: Correct, correct. So nothing will change from the user standpoint. But there'll be much tighter authentication using OAuth 2.0 for the apps. And they'll be bringing in, they'll be shortly bringing online an SDK, a Software Development Kit, for apps, telling them we're lowering the boom on you on September 1st of 2011. So you have between - I think the SDK is supposed to be available in June or July timeframe, so a couple months from now. But all - and it's already running. I mean, Facebook themselves has switched over.

So they're making good improvements. And all apps will have to switch to that. In the meantime, they are no longer leaking this user token in the URL that they provide to apps that are wanting to authenticate and then act on behalf of users. But all the old ones do not expire. And Symantec has said the only way to fix it is for the user to change their password to invalidate this. So my sense is it's broad, and it's distributed. But, I mean, change your password.

Leo: I can't believe that this isn't, like, I mean, this is, well, okay. Unless I'm misunderstanding what they can do with this. I mean, they can't - all they can do is post to your wall, stuff like that. Whatever permissions you gave these apps.

Steve: Yes. That token is absolutely restricted to whatever it is you gave the app permission to do.

Leo: Although this is what the bad guys do. They post something on your wall that says, hey, you've got to see this video of the new TWiT cottage, and it links to what looks like YouTube, and then it says, oh, your Flash is out of date, and you say, oh, I'd better update it. And then you've got malware. So that's really - it's things like that. I mean, I don't - my profile is public, so I'm not worried about that. My son is.

Steve: As you said, Leo, it is people, for example, going to your Facebook page, that your page then infects.

Leo: Right.

Steve: And I didn't...

Leo: Could it do that? Could it do that, a spontaneous infection?

Steve: Oh, yeah.

Leo: Oh, crap.

Steve: I was just going to say, I didn't - we talked last week about Google Images, the problem that we're seeing with Google Images where the images are malicious.

Leo: Right. Injection.

Steve: Yes.

Leo: And that's not coming from Google, but on the image search you get an image from the offending page which has an infection in it.

Steve: Exactly. And I did see reference in some forums in the last week where people were commenting about the problem with Google Images and that just in hovering over the image where it zoomed in, like magnified it, that grabbed his computer.

Leo: Fudge.

Steve: I mean, that was...

Leo: And that's the kind of thing of course Facebook, you know, it's just a natural.

Steve: Yes. And so here's the problem, is that the scenario would be the bad guys who now know about this will get the server logs of, like, advertisers or anybody with an affiliated relationship to these apps, and now they know to scan the referrer field for these tokens. And those tokens will allow them to act as if they were the app to which the user had given permission. Now, you could revoke the app permission, and that would protect you, too. But it's easier just to change your Facebook password.

Leo: Yeah, especially since you probably have hundreds of app permissions, as most people do, I mean...

Steve: Yeah, yeah.

Leo: I'm just, again, I'm stunned that this isn't, I mean, holy cow.

Steve: I don't have a date here. I don't know how old this is. But this is just - it just happened.

Leo: Well, I'm changing my password right now. And why not, anyway. It's a good thing to do from time to time anyway.

Steve: Yeah. The title was "Facebook Applications Accidentally Leaking Access to Third

Parties." And, whoops. Oh, and that's actually the title of Symantec's blog.

Leo: It came out yesterday.

Steve: Okay. So it's been just recent.

Leo: There should be a big banner on Facebook saying...

Steve: Please, everybody...

Leo: ...change your password.

Steve: ...change your password. Yeah, you're right. They really ought to step up and take responsibility for this. Because, I mean, otherwise it's going to be - it'll never get enough attention. Users won't change their passwords. And there will be diffuse attacks. There will be, where these logs are available, bad guys will pillage them and sort through them, find the tokens, and then get up to mischief.

Leo: I'm trying to find - of course you can't find anything on Facebook. I'm just trying to find in my account settings where I've given authorization. Let's see, apps and websites. So you have to go to Account - this is ridiculous - Account > Privacy Settings > Apps & Websites. Down at the bottom, of course, not where you'd think it is. I just went through this with my son last night. He said, "People can see my pictures?" I said, "Well, yeah. You didn't know that?" He said, "No," because the default is they can see his pictures. So I showed him how to change that so only friends could see his pictures. Now I've got to say, "And by the way, Henry, change your passwords."

Steve: Just "password." Just your master login password.

Leo: Just your password.

Steve: Yeah.

Leo: So Pulse, Eventbrite, Empire Avenue, Seesmic Web, Foursquare. No, there is one button, turn off all platform apps. But then I'd have to change all of them. I'd have to log them back in. So it's easier just to change a password.

Steve: Yeah.

Leo: Crazy.

Steve: Yes. And speaking of getting up to no good, I just thought, oh, I wonder how many Firesheep downloads we've had so far.

Leo: This is Firesheep-like, isn't it. Kind of the same result.

Steve: Well, yeah. I mean, it is an impersonation attack. The bad news is, this has been happening since 2007, Leo. So, I mean, ever since Facebook began making apps available, if the app was hosted in an iFrame, then this iFrame enabled the leakage. And then if the app ever pulled any third-party content, the referrer field leaked that URL that the app used to the third party. And that is a static persistent token that allows that third party the same access privileges as you had given the app. So, I mean, for all time, until, like, a few days ago, when Symantec said, uh, gee, Mark, you've got a problem over here.

Leo: Wow. I'm just...

Steve: But anyway...

Leo: I'm just, like, stunned by this.

Steve: Yeah. We are approaching 1.5 million downloads of Firesheep. Last time I looked we were at 1.3. And we are at 1,492,829 when I last looked, although every time I looked, like an hour before that it was at 549. So about 300 per hour, seems. And so we'll cross 1.5 million downloads of Firesheep here in another week or two, probably. Now, this is...

Leo: I'm just depressed. I am just so depressed. Because nobody - this is massive, and nobody's going to change their Facebook password. My wife isn't. My kids aren't. Their friends aren't. Nobody's going to do this.

Steve: Yeah, I wonder that Facebook couldn't obsolete those tokens.

Leo: Of course, of course they could. But they won't because they know they would get a hundred million phone calls. Can you imagine the cost? When you have...

Steve: Well, it's like LastPass last week. The site was down for a couple of days because they told everybody they had to change their password.

Leo: And they probably only have a few hundred thousand users. And by the way, they did exactly the right thing. They didn't even have evidence, they didn't even know that there was an infraction. They just through there might be.

Steve: Yes.

Leo: And Facebook, knowing this, just blithely goes along. I wish I could delete my Facebook account. I can't because it's what I...

Steve: Well, you can just change your password, though.

Leo: Well, I know. But this is the tip of the iceberg.

Steve: Oh, I see. I see. You're again upset enough over this that - yeah.

Leo: Well, and it's the tip of the iceberg. We don't know what other crap's going on on Facebook. We will never know. I mean, obviously they don't want to tell you. Wow.

Steve: Yeah.

Leo: Okay, sorry.

Steve: Maybe it's, I mean, in the same way that our financial institutions were too big to fail, maybe Facebook now is too big to tell everyone to change their password. It would just bring down the Internet.

Leo: It would. Oy gevalt. All right. All right.

Steve: Okay. So Juniper Networks, that's not a security company, they're a big iron router manufacturer, essentially, Juniper Networks has been around forever. They produced a report which I'm just going to - I've pulled some tidbits from. And they called it their "Mobile Threats Report for 2010 and 2011." So this is a little historical, reaching back, looking at what they've seen before.

So the key findings of the report, they called it - one is "App Store Anxiety," is that "The single greatest distribution point for mobile malware is application download, yet the vast majority of smartphone users are not employing an antivirus solution on their mobile device to scan for malware." I mean, we don't really have that yet.

"WiFi Worries: Mobile devices are increasingly susceptible to WiFi attacks, including applications that enable an attacker to easily log into victim email and social networking applications."

"The Text Threat: 17 percent of all reported infections were due to SMS trojans that sent SMS messages to premium rate numbers, often at irretrievable cost to the user or enterprise."

Leo: Oh, yeah, that's a hell of a scam, yeah.

Steve: Yeah. "Device Loss and Theft: One in 20 Juniper customer devices were lost or stolen, requiring locate, lock, or wipe commands to be issued." So 5 percent. "Risky Teen Behavior: 20 percent..."

Leo: Using Facebook.

Steve: "20 percent of all teams admit sending inappropriate or explicit material from a mobile device."

Leo: Oh, everybody does that.

Steve: And then, finally - those darn cameras. And then finally, "Droid Distress: The number of Android malware attacks increased 400 percent since Summer of 2010."

Leo: That's - to me, a percent, I don't want to know that. That's, what, that's four times - if there was one, that's four. I want to know the number. Give me the raw number.

Steve: Yeah, good point. And then, quoting from the report, they said: "These findings reflect a perfect storm of users who are either uneducated on or disinterested in security, downloading readily available applications from unknown and unvetted sources in the complete absence of mobile device security solutions." And this is Dan Hoffman, their chief of mobile security at Juniper Networks. He said, "App store processes of reactively removing applications identified as malicious after they have been installed by thousands of users is insufficient as a means to control malware proliferation. There are specific steps users must take to mitigate mobile attacks. Both enterprises and consumers alike need to be aware of the growing risks associated with the convenience of having the Internet in the palm of your hand."

Now, their suggested actions are not going to impress any of us. They say "Install an [add-on] antimalware solution to protect against malicious applications and spyware, infected SD cards, and malware-based attacks on the device. Use an [add-on] personal firewall to protect device interfaces. Require robust password protection for device access. Implement antispam software to protect against unwanted voice and SMS/MMS communications. For parents, use device usage monitoring software to oversee and control pre-adult mobile device usage and protect against cyberbullying, cyberstalking, exploitative and inappropriate usage and other threats."

So when I stand back and look at all of that, what I see, Leo, is an immature piece of the industry. It's that phones are new. And users are new to this kind of phone. PCs are much more mature from a security standpoint. Now all personal computers have firewalls built in. And I guess there's a higher bar to using a PC securely than there is a phone. A phone just sort of - it's hard to take it that seriously. Is that what you think it is?

Leo: Yeah. Yeah. And it's also hard to secure it. So it's not obvious.

Steve: It is. I mean, we know, for example, that the threat that PC users have is going

to malicious web pages. Sort of the similar threat that is still, I think, unappreciated is just malicious free toys for these phones. I mean, I have hundreds.

Leo: I download stuff all the time. And who's checking this stuff, you know?

Steve: Yeah. I mean, I see something interesting for the iPad, it's like, ooh, there's a neat toy.

Leo: We presume that Apple is vetting these. But the problem, as you can see, is it's almost impossible to be, you can never be a hundred percent sure that something is secure.

Steve: Well, and the reason Android is popular is that it doesn't have the boot of Apple on it to the same degree. Unfortunately, what that also means is it is a larger target than Apple is. And, I mean, Android, I'm sure you're seeing the numbers.

Leo: Oh, yeah. If I'm a bad guy, I don't worry about the iPhone. I go right after Android, absolutely. Yeah, much easier. It's the Windows of the phone world.

Steve: Yeah, it is. Okay. Now, in another - this is, I guess, not quite my last bit of good news for the day. We have a new attack vector which is presenting itself known as WebGL. This is the next generation of web-based 3D graphics. And essentially, OpenGL has been around for years. It's a sophisticated, mature API, an Application Programming Interface, to allow applications access to powerful, hopefully powerful rendering hardware in the machine. The problem is that, in the same way that there's nothing inherently wrong with scripting unless you go to a malicious site that scripts you maliciously, well, there's nothing wrong with OpenGL unless you go to a malicious site that uses that technology through what's known as WebGL, maliciously. And it turns out it's possible. All of the latest browsers are now supporting WebGL.

Leo: How interesting.

Steve: Yes. And listen to this. The technology of this is interesting. I wanted to just quote from Context Info Security's site. They said: "Traditional browser content would not normally have direct access to the hardware in any form."

Leo: No, it seems silly on the face of it.

Steve: Of course. "If you drew a bitmap, it would be handled by some code in the browser with responsibility for drawing bitmaps. This would then be likely to delegate that responsibility to an OS component, which would perform the drawing itself." Which is the way our stuff works.

"While this distinction is blurring somewhat with the introduction of 2D graphics acceleration in all the popular web browsers, it is still the case that the actual

functionality of the GPU" - the Graphics Processing Unit - "is not directly exposed to a web page. The salient facts are that the content is pretty easy to verify, has a measurable rendering time relative to the content, and generally contains little programmable functionality." And that's the key. An image is just that, it's an image.

"WebGL, on the other hand, provides by virtue of its functional requirements" - so this is not a mistake, this is the way it was designed - "access to the graphics hardware. Shader code, while not written in the native language of the GPU, is compiled, uploaded, and then executed on the graphics hardware. Render times for medium to complex geometry can be difficult to determine ahead of time from the raw data, as it is hard to generate an accurate value without first rendering it - a classic chicken-and-egg issue. Also, some data can be hard to verify, and security restrictions can be difficult to enforce once out of the control of the WebGL implementation.

"This might not be such an issue, except for the fact that the current hardware and graphics pipeline implementations are not designed to be preemptible or maintain security boundaries. Once a display list has been placed on the GPU by the scheduler, it can be difficult to stop it, at least without causing obvious, system-wide visual corruption and instabilities. By carefully crafting content, it is possible to seriously impact the OS's ability to draw the user interface, or worse. The difficulty in verifying all content and maintain security boundaries also have potential impact on the integrity of the system and user data."

So what these guys are saying, and they have done proof of concept, they have been able to blue screen people's machines by visiting a web page with maliciously crafted 3D graphics. And the problem is that, in order to do what - as we know, GPUs are very powerful. Essentially the web server is loading code into your GPU, which it then runs. And their recommendation, and actually US-CERT has recommended looking at the page that I provided a link to here and disabling WebGL if it is present on your browser.

Leo: Wow, that bad.

Steve: It is that exploitable, apparently.

Leo: We don't need to say it on this show, but just to reiterate, a blue screen is always the first step to pwning a computer. If you can crash it, then you just have to figure out where to write the code. It's not so hard.

Steve: Exactly. It's a matter of finesse after you've demonstrated that, oops, you're able to kill the...

Leo: That's always, yeah, always the first step.

Steve: So apparently the latest versions of Firefox, Chrome, and Safari all support WebGL.

Leo: Oh, yeah.

Steve: And Opera has just released Opera 11, a preview, that supports it.

Leo: I've used it. I mean, it's cool. But I guess it's not so good.

Steve: It's cool. And unfortunately, as most cool things, there's a dark side.

Leo: There's a dark side. We continue on with Attacks & Breaches.

Steve: So I thought probably that's where this should go, although this wasn't a website that was attacked and breached, it was the famous Google Chrome sandbox that was breached. Some developers using two previously unknown zero-day vulnerabilities were able to break out of the Google Chrome, the latest version of the Google Chrome browser sandbox and then not only get out of the sandbox, but then circumvent address space layout randomization, ASLR, and the data execution prevention, DEP, the two major technologies in the latest versions of Windows that are now on and active and used by Google in order to make sure that the components that are in their process space are in random locations and that you're not able to execute data. And they were able to, using an exploit, essentially just going to a web server, able to download Calculator.exe from somewhere else and run it at minimum integrity level. So, I mean, that's all you need to run, I mean, their demo downloads Calculator.exe, but it could also be malware like you've never seen before dot exe.

Now, what's really weird is I don't know what these guys are up to, exactly. This is VUPEN.com, "pen" as in penetration testing. And if you just go to VUPEN.com, down sort of like the top item in the lower part of the page, they'll talk about pwning Chrome. And Kelly Jackson Higgins, who is writing for Dark Reading, sort of summed things up. She said, "VUPEN - which withheld technical details of the bugs in its disclosure - had not disclosed the bugs or any details to Google as of this posting."

Leo: Oh, I don't - that's not good. I don't like it when they do that.

Steve: I know. "The security firm provides details of vulnerabilities it discovers to its paying government customers. Quote, 'We did not publicly disclose any technical details of the vulnerabilities for security reasons." Well, good, publicly, of course. "We did not..."

Leo: That's fine. But tell Google.

Steve: Yes. "'We did not send the technical details of the vulnerabilities to Google, and Google did not ask us to provide these details,' says Chaouki Bekrar, CEO and head of research at VUPEN." And then what of Google? A Google spokesperson said in a statement that without any details of the hack, the company is unable to verify it. So, quote, "We're unable to verify VUPEN's claims at this time as we have not received any details from them. Should any modifications become necessary, users will be automatically updated to the latest version of Chrome." So said the spokesman. And what VUPEN is doing by way of proof is on their site and on many sites that have picked up this story about the Google Chrome sandbox being pwned. They have a YouTube video showing them doing the exploit, basically going to a page, and the act of going

there causes Calculator.exe to be launched, and that should absolutely be impossible because...

Leo: Yeah, but it's a video that they, I mean, it's going to a local URL, I mean, it's not really a proof of concept. It's just a video.

Steve: Right. Well, no...

Leo: It's meaningless.

Steve: Precisely. So the problem is...

Leo: It's not proof.

Steve: I don't get what it is that they're - I don't understand what it is that they're - what game they're playing.

Leo: Well, they're saying we share these with our government customers. So we aren't going to tell you because you don't pay for it.

Steve: Yeah.

Leo: Which sucks.

Steve: Yeah. I mean, Google will pay them for these - they've got two zero-day vulnerabilities. One gets out of the sandbox, and the second one, once out, is then able to do the work outside in order to get an arbitrary executable downloaded from anywhere on the Internet. I mean, those are powerful. And so one hopes that the government, whoever it is that they're selling these things to, are being responsible with them. But I don't...

Leo: I don't know.

Steve: Just scratch my head on that one.

Leo: Yeah, I don't - that's not - this is not good behavior.

Steve: It does not seem like the right thing to do. I did want to follow up on, from one of our listeners, a D.M. Ovad, who said - who just sent email actually to my company saying please pass this on to Steve. He said, "Steve, from last week's Security Now! podcast, I believe Episode 299" - which, yup, was last week - "you mentioned the

possible high temperatures while using SpinRite on a laptop. A few years ago, while preparing a hard drive on my wife's Fujitsu laptop, SpinRite did warn about the elevated temperatures of the drive and paused its processes. All I did was to fill a one-gallon Ziploc bag with ice cubes, laid it on my desk, placed a washcloth over the bag, then the laptop on the washcloth. Of course I positioned the laptop so the location of the hard drive was directly over the ice pack. SpinRite continued, and the hard drive's temperature remained well below the maximum, and SpinRite was able to complete and successfully repair my wife's hard drive."

Leo: Don't you worry about condensation, though, with something like that?

Steve: Well, I think that's where the washcloth comes in. So he had a nice soft insulating layer and managed to keep his drive cool. He says, "Thanks for a great solution to many hard drive issues. Dave." Oh, so his first name is Dave.

Leo: I love it. Thank you, Dave.

Steve: D.M. Ovad. So, yeah. Thank you for the tip. I thought I'd pass it on to any listeners who may have overheating hard drives and need to run SpinRite anyway.

Leo: Now, this is a Q&A episode, but we're about an hour and a half into it. So we've got some tweets questions for you.

Steve: We do. I did want to quickly mention that OpenDNS, that we have spoken of often...

Leo: Yes, we have.

Steve: ...is now supporting IPv6.

Leo: Oh, wow, that's great.

Steve: In advance support for World IPv6 Day, which is June 8th and approaching, that's next month, about a month from now, they are now supporting IPv6. And in the announcement they said that their IPv6 addresses for the OpenDNS IPv6, what they called a "DNS Sandbox," which is to say the IPv6-enabled DNS servers are 2620:0:ccc::2 and 2620:0:ccd::2.

Leo: Wow, that's interesting. I'd never seen IPv6 addresses.

Steve: I've just - that's why I wanted to state them and show them, as our listeners will certainly be, in the future, hearing IPv6 addresses.

Leo: What is the - why is it "::" at the end?

Steve: The way this works is, and this has been well thought out because the problem, of course, is that 128 bits is four times longer than the 32 that we've been using. The 32 that we've been using used the so-called "dotted quad" format. Well, we would have to have 16 of those if we didn't come up with a way of compressing it. We'd have to have, like, 16 numbers ranging between zero and 255. So instead what they do is they set them up as groups of 16 bits, representing by four hex characters. So first of all, we're no longer decimal, we're now in hex.

Leo: Hence cc, cc, ccc...

Steve: Exactly. So ccc. And there's always an - now, I said ":ccc:", there was only three c's. But hex for 16 bits needs four. So there's always an implied leading zero, if it's not specified. And the "::" means that there are as many zeroes in between the two colons as necessary to push each side out to the ends.

Leo: Ah. So it's like a fill colon.

Steve: Yes, exactly. So, and the way IP space is allocated, you'll generally have, like, a left-hand-side network and then a right-hand-side machine, or subnet. So that's what we're seeing, where it's 2620:0:ccc, and the other one is 2620:0:ccd. And then in both cases they end with a "::2", meaning then just do all zeroes until you get down to the end, and then we have a 2 at the very end. And 2 will be a common number for machines and interfaces and so forth.

And I wanted to quickly say something to our listeners that I tweeted. You probably know about this already, Leo. But I went to IMDB on my iPad, the Internet Movie Database? And it said "We have an app." There's an app for that.

Leo: Oh, yeah. There's a great app.

Steve: And it's lovely. And that's the end of my message.

Leo: I could show it to you. I have my iPad here. But I also have my - I have an Android tablet, so I could show it to you, if you want to see it I could show it to you on the Android tablet here.

Steve: Anybody who is - yes. So it's iPhone, iPad, iPod Touch, any iOS device, or Android. And if you're a movie person, and you and I are, Leo, and I'm sure a lot of our listeners are, it's just - they just nailed it. It is simple and elegant. You're able to just, like, I love that they deal with, like, movies coming up in the future. You can scroll through them. You touch one, you go there. You touch an actor, you go to the actor's profile. It just, like, interlinks all of that so nicely. It's just...

Leo: I thought I had it, but I don't see it here. So I'll - but I do, yeah, I love it. I agree.

Steve: And I guess there are other IMDB apps. Other people tweeted me...

Leo: Yeah, but get the official.

Steve: The official one they just nailed it. They did a great job. In tweets from the field, Von Welch, whose @VonWelch is his handle, he mentioned covering webcam on MacBook Pro with sticky...

Leo: Hey, before you do that, can I do an ad?

Steve: Oh, yeah.

Leo: I'm sorry. I haven't been paying close attention. And here I am trying to log into Facebook.

Steve: Well, so Von Welch noted, he said, "Cover webcam on MacBook Pro with sticky, Mac thinks I'm in dark room and dims the screen."

Leo: Oh, that's interesting.

Steve: Turn off auto-dimming under Preferences/Display.

Leo: Oh, that's a good point.

Steve: So, yes. Apparently the MacBook uses the camera also as its ambient light sensor for controlling the automatic screen brightness. And I was pleased that Von Welch decided to follow up on my advice after hearing about another instance of people being spied on with their webcams. But when he did it, his screen got dim. So you can turn off auto-dimming under the Display panel under Mac Preferences. So I just wanted to pass that on to our listeners.

And this is very cool. Graham Wetzler tweeted about a security tool or I guess service that I really like. It's the kind of thing I wish I could have done, or I have done. But now it's been done, so I don't need to: urlxray.com. Give it a shortened URL, and it'll tell you where it leads.

Leo: Oh, that's good.

Steve: Which is trivial to do. I mean, all it does is it goes to the URL. And the way these redirectors work is they return you a 301 HTTP Moved response with a location header telling the browser where it should jump to. So, I mean, the trick is simple. But obviously instead of jumping to it, it just shows you where your browser would jump to, if you had given that URL to it. And I have to say, I mean, TweetDeck does a good job at showing those often. But when people tweet me something and allow me to sort of get a sense for them and, like, what this link is they're providing, I do have a secure machine. But I'll typically trust people and click it. But sometimes I get a tweet that just says, "Oh, check this out." And it's like, I'm not clicking that, honey.

And it reminds me of the famous scene from "The Wrath of Khan," where somebody is telling Kirk to raise his shields because he's not hearing anything from the oncoming starship. And of course it's Khan, who then blows him up because he didn't have his shields raised. And the idea that, it's strange, but if you get some communication from someone, you have some sense for who they are and what they're about and where it might be that you're sending your browser when you click the link. But it's very different if you just get a URL and nothing. It's like, well, okay, I'm just - I'm reluctant to click that. So anyway, urlxray.com allows you to disambiguate these URLs prior to going there directly, which I think is very useful. I'm going to be using it. In fact, I may just implement it myself for GRC, and then we'll have one more little gizmo at GRC.

Okay, so real quickly, a couple of Q&As. Niko Carpenter tweeted, and several people had asked who sent feedback to GRC.com/feedback, our regular place, he said, in a split handshake, how does a bad guy get the server to only send a SYN packet? And I loved that because I completely forgot to talk about it. Remember? Okay. So split handshake was - I talked about a week or two ago where it was a potential exploit where intrusion detection software could be debilitated, essentially confused by the exact sequencing of TCP packets back and forth. So normally you send a SYN packet to the server, it returns a SYN ACK, and then you ACK it, and now you're connected, and you proceed. In a split handshake, instead the server sort of ignores, basically ignores your SYN and just sends a SYN to you, essentially sort of turning the handshake around. And that's what confuses the IDS, the Intrusion Detection Systems. So he says, how does a bad guy get the server to send only a SYN packet?

Well, what I forgot to explain was that you would have to be going to a bad server in the first place. So it's an exploit that the server, that only the server you are reaching out to can perpetrate against you. So it's not like bad guys somewhere can use this to, like, get Google to connect to you, but you're really connecting to them or something. It's, you know, if you were going to a site that wanted, that already wanted to do you harm, then that site would send a SYN packet back to you, rather than a SYN ACK, reversing this connection and using that in order to bypass intrusion detection software that would otherwise be protecting you. So that's how that works.

Luis Fernando asks, whatever happened to that capacitor idea you announced a while ago?

Leo: Oh, the supercapacitors, yeah.

Steve: The supercapacitor. And it's funny because I had myself looked and wondered what was going on. And this was a company called EEStor down in Texas. And last we heard from them, they were in the process of working out the details of mass production. And they disappeared.

Leo: Oh, dear.

Steve: My sense is they probably could not make them, whether it was a quality control problem, a mass production problem, a couldn't hook enough of them together problem, we don't know what.

Leo: We know they're for real because I actually have a screwdriver that charges superfast.

Steve: Yes. No, you mean supercaps are real.

Leo: Supercaps, yes. We just...

Steve: Yes, supercaps are definitely real. And in fact...

Leo: These guys were going to make it for cars.

Steve: Yes. And there was an automotive manufacturer in Canada that was getting all geared up and ready to go. Interestingly, I saw a quote from the CEO of Tesla. And we all know Tesla, the manufacturer of that fantastic little sports car that is ridiculously fast. One and I were next to each other at a stop sign a couple weeks ago, and this thing just shot off, like, when the light turned green. And I thought - and, I mean, silently. Didn't make any sound. I thought - and I looked, and sure enough "Tesla" was across the back of it.

Leo: Awesome, yeah.

Steve: The CEO stated, "Supercapacitors are the future of automotive travel." This is the CEO of Tesla said, yes, we're only doing batteries now because we don't yet have supercapacitors. But remember, I mean, the reason I was so bullish about them is that they solved the wear-out problem, and they solved the slow charging problem. If you had, like, a high-current pumping station with high-current delivery to the car, and the pumping station probably had supercaps, too, so it was always sucking power off the grid, filling up its own supercapacitors, then you bring your car in, plug it in with some mondo connector, and, I mean, many, many, many hundreds of amps, and the charging station would dump its supercaps into your supercaps, and off you'd go.

So I'm convinced it's the right technology. And the CEO of Tesla, who's got a lot of experience with this, agrees with me. We just don't have them yet. But many people, I mean, I'm seeing work happening. There's, like, there's lots of university research with nanotubes and all kinds of strange things that are, like, working on getting supercap technology to happen. So I think we're not far from it.

Leo: Good, good.

Steve: And finally, Carlos Cardona also twitted. He said, you often talk about password hashing/salting. Could you talk about password stretching? And I touched on it last week when we talked about LastPass because that is the - we understand that, when you hash a password, you pass the plaintext through this cryptographic hash, and it turns it into a fixed-size token that represents what you fed in. We know that when you salt the hash, you add something secret to the password, and then you hash them both. And the benefit of that is that you then don't - because hashes use standard formulas, like an SHA-1 or an SHA-256, or MD5 is no longer considered very strong. But the hashing algorithm itself is a universal standard.

So the problem is, if you didn't add salt, then somebody could figure out your password by brute forcing all possible, like, reasonable passwords through that hash to see if they get the same token that was stored, for example, when a secure website had stored the token. But by adding hash, I mean, sorry, by adding salt, you essentially create a custom hash function. That's what the salt does is it sort of customizes the hash function so that, by mixing in with the input you're providing, you're going to get a token out different than that universal function would otherwise provide if it weren't salted. So that we have.

The one problem we still have is what if somebody got the salt? And that was the concern, unverified, and still unverified by the guys at LastPass. They were concerned that somebody may have gotten the salt, which would then have weakened the protection salt provides.

So there's one more thing we can do, and that's called "password stretching." And what it does is it stretches it in time rather than in size. And it's as simple as, rather than hashing once, you hash many times. For example, the WiFi folks who did WPA, the current standard good secure WiFi, they understood this. So they salt the hash with the access points, access point name, the ID of the access point, so that when you set your access point's name, you are increasing your security by mixing that salt in. And then they repeat that 4,096 times, the idea being that anybody brute forcing can't just do it once, they've got to do it 4,096 times, which slows down their brute-forcing speed by 4,096. No matter how fast they can do the hash, they're going to have to do it 4,096 times per brute-force attempt.

Leo: Even if they have the database they still have to go through this.

Steve: Yes. Yes. And that's the beauty. And, for example, in the LastPass case, they're talking about doing it 100,000 times.

Leo: That should be enough.

Steve: So it'll really slow it down, yeah.

Leo: Wow. So this has been a great show. And I think if nothing else I hope people have learned a little something from this Facebook thing. This is why you listen to

this show, folks. Because I don't know, I can't understand why this isn't being more publicized. PC Magazine just published an article about it, and I've linked that on my Twitter account. Wow.

Steve: Yeah. Well, I imagine it'll get picked up by...

Leo: I hope so.

Steve: ...Techmeme and Boy Genius and those guys, and then hopefully then by the mainstream. And ultimately you have to imagine that Facebook will be forced to do something.

Leo: Well, they've already said, oh, yeah, it's not a big deal.

Steve: Okay. Yeah. As I said, my sense is it will allow diffuse attacks, meaning that because it's a function of that authentication token being found, and that only allows them access to one Facebook user's account. So there's a lot of Facebook users. And...

Leo: Well, somebody would really most likely create a malicious app to take advantage of this, if you were going to do this. Facebook's response, they said "We've worked with Symantec to rectify the issue," but took issue with how Symantec characterized the situation. "We've conducted a thorough investigation," says Facebook, "which revealed no evidence of this issue resulting in a user's private information being shared with unauthorized third parties. In addition, this report ignores the contractual..."

Steve: Oh.

Leo: Oh, dear. I'm starting to - "the contractual obligation of advertisers and developers which prohibits them from obtaining or sharing user information in a way that violates our policies."

Steve: Oh, goodness. Well, and Part A there is, well, we weren't looking to see if this has been happening, so we didn't see it happening.

Leo: There's no evidence. Well, how would you know? What, are you going to look back at all the logs for - it's been going on for years.

Steve: And people's Facebook accounts are being hacked all the time. Gee, I wonder how.

Leo: Symantec says, "We estimate that over the years hundreds of thousands of applications may have inadvertently leaked millions of access tokens to third parties.

Steve: Yes.

Leo: But we have no evidence, says Facebook. And by the way, it would be a violation of their terms.

Steve: Oh, they didn't read the fine print.

Leo: Oh, I'm just - I am just done with Facebook. It's so - I'm so done. Holy cow. Just horrible. Steve Gibson is at GRC.com. There at the bottom of the screen are all his Twitter handles: @SGgrc is his main Twitter account. @SGpad for pad for tablet-related news. And @GibsonResearch, the official account for his company, the Gibson Research Corporation. Of course that's GRC.com. That's the place to go if you want to get SpinRite, the world's finest hard drive maintenance and recovery utility. For questions in future episodes, GRC.com/feedback is a good place to go. I also suggest that you check out his other free programs and all sorts of information at GRC.com.

Steve: Okay. I'm going to do something really strange, Leo.

Leo: Okay, I'm ready.

Steve: My Passcode Designer that I've talked about several times came alive yesterday.

Leo: Ooh.

Steve: And it is - I think it's going to live up to my expectations. Probably by next week, I'm going to tweak the graphics a little bit - and this is my big JavaScript program. This is my project. I taught myself JavaScript so I could create this thing because it runs on a web page. This idea just hit me because the guys in the GRC user group are really anxious to see it. I've been talking about it more to them than I have been on the podcast. But I'm going to introduce it next week to the listeners of this podcast as something of a puzzle.

Leo: Fun.

Steve: Because it will exist, but the documentation won't. And it's a little machine that I've built. And when you click things and do things, things happen. And I realized people could scratch their head and think, what is this doing?

Leo: Might be kind of fun.

Steve: I mean, you poke it over here, and it does this; and you poke it over there, and that happens. And so next week I will - I'm sure I'll be ready by then because it's running now. I'm going to give everyone access to it and just say, I haven't written anything yet. There's no documentation. But if you want a toy to play with that involves passcodes, you can poke at this and see if you can figure it out because it's something that needs to be figured out. So...

Leo: Well, don't tell anybody. We'll find out next week.

Steve: Yup. And we will - we're going to plow into randomness, how we solve the problem of needing random numbers from a computer that cannot make them. It can't. That's not what it's for.

Leo: Part 2 of Going Random. Steve Gibson.

Steve: Thanks, Leo.

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