SECURITY NOW!

Transcript of Episode #330

Listener Feedback #132

Description: Steve and Tom 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-330.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-330-lq.mp3

TOM MERRITT: This is Security Now! with Steve Gibson, Episode 330, recorded December 7th, 2011: Your questions, Steve's answers, #132.

It's time for Security Now!, the show you need to watch if you want to stay safe on the Internet because of this man, the guy from GRC.com, SpinRite, ShieldsUP, all kinds of good stuff, Mr. Steve Gibson, with us today as usual. But I'm with you today as not usual. Leo Laporte's out in France. Pleasure to be with you again on Security Now!, Steve.

Steve Gibson: Likewise, Tom. We always have a good time, and you've got the whole beat down for the podcast, so - and somehow it always seems you land on the even episodes, which are Q&A episodes.

TOM: I love the Q&A episodes.

Steve: I mean, you've done non-Q&A, I think, once or twice. But generally it happens that you're on the Q&A, so...

TOM: Yeah, I lucked out. I got to be on one of the How the Internet Works episodes, which was fantastic. That was a great tutorial series. And I like the Q&A episodes, too, because they're little samples of lots of different stuff.

Steve: It's funny because I'm - I guess I'm really more conscious of the continuity that we have from one episode to the next. And, like, for example, there's some things here that I'm aware that I would like Leo to know.

TOM: Oh, yeah.

Steve: Because it'd be nice to have that for the future. And it's like, well, okay, Tom's going to know this, but Leo's not going to know this. So our listeners all know it. But it helps when the person I'm working with also knows it. And I'm thinking, oh, shoot. Well, Leo's going to miss this. But I'll sort of try to remember that he doesn't know that, or he'll say something.

TOM: Or he could watch the show.

Steve: If he's listening, yeah. He's not, though.

TOM: He might be listening right now, going "I'm hearing you, Steve. I'm getting it all."

Steve: I don't think so. He's in France. He's not listening to the Security Now! podcast.

TOM: Probably not. He's probably asleep. But he can listen to it later on the podcast. Let's get into the news, Steve. You always seem to start with a zero-day something. We have a zero-day flaw in Adobe Reader this week.

Steve: Adobe gives us so much to work with.

TOM: Yeah, they do, don't they.

Steve: Yeah, my goodness. So the big news is that Adobe alerted the security followers, the security industry, to their detection of a zero-day, a new zero-day flaw in Adobe's Reader and Acrobat products. So what they're seeing is they're seeing this, because it's zero-day, meaning the first time they learned about it was someone reported that this was actually being exploited. So there are flaws in, not only the current version, which is 10.1.1, and all earlier. So these are being used in targeted attacks where email is being sent containing a PDF which is deliberately malformed in order to take advantage of a weakness in the security in the handling of some aspect of PDF documents.

Now, what's significant is that the new v10 of Reader and Acrobat has the so-called "Protected Mode," or in Acrobat I think it's called "Protected View." And in this case those new protected modes have mitigated the problem. The problem exists in these versions, but it can't get out. It cannot do the damage that it can in the pre-v10 releases. So...

TOM: So you still want to update v10 to get rid of this.

Steve: Well, actually Adobe's not even going to make it available.

TOM: No, they're not. Okay. Well, then they're not.

Steve: Yeah. They're that confident that this does - that their containment technology does keep it corralled. No, that was the perfect question to ask. So what Adobe's going to do is, they will be updating to fix this problem in the v10 line, but they're not going to do it in an emergency fashion. It'll be their regular quarterly update schedule.

TOM: Gotcha.

Steve: Which, for a change, they're going to stick with. They haven't been ever sticking with it until now. But in this case, this exploit, which is not to say that all exploits, but this one does get contained by their containment technology that they introduced in v10. What they are going to do, though, is an emergency release, which they have committed to have available for all v9 and prior versions of Reader and Acrobat, no later than the week of December 12th. So, what, we're in the week of December 5th. So that is to say, so no longer than next week, the week following this podcast, there will be an emergency release from Adobe. They have promised to fix this problem because it's a bad one for people who did not move up to 10.x of Reader and Acrobat.

TOM: It's another good example of why you should maintain updates and always use the latest version.

Steve: Yeah. And there are, largely in corporate settings, there's more inertia because they're feeling like, well, we already know what v9 does, and it's working. And there's my - I'm famous for my reticence to move to the latest and greatest because oftentimes you just end up with some different shaped arrows in your back. I'm still using XP. And I look over here, and I see that we have 852 days left of support for Windows XP. So I'm not in a big hurry.

Although I have seen some things in the news that talked about - oh, in fact it was a little blurb that didn't quite make my notes for this week because it wasn't anything spectacular, but it mentions that - and I was wondering what the source of this was. It's probably Windows, I mean Microsoft, because there was some mention that, had RSA been using Windows 7 rather than Windows XP, the security inherent in Windows 7 would have prevented the exploitation of the flaw, which is what compromised the RSA SecurID keys which we covered extensively some time ago. So it's like, okay, well, that's good for them.

But we do know that Windows XP is still the majority platform on the Internet. That is, Microsoft is not, in general, succeeding in getting people to move off of XP because there are old stick-in-the-muds like me for whom XP is still doing everything I need. And it's nice to let a few service packs get out there on a new OS before you jump to it. Which was the case with Windows 7. There were new security problems with Windows 7 which have now been resolved.

TOM: But suffice to say, when there's an update for security, make sure you implement.

Steve: Yes, yes, yes.

TOM: That's what I was trying to get at.

Steve: Yes. And, yeah, absolutely. And I don't - for our listeners, typical end users, or maybe people in corporations who can work to drag their companies forward, certainly being up at Adobe Reader 10, I mean, it's free. Why wouldn't you be using 10 versus 9?

TOM: Yes, unless you're using Foxit or something.

Steve: And don't have the problem at all.

TOM: Yeah, exactly. All right. Let's move on to the OpenDNS "DNSCrypt" beta for Mac users. I thought this was really interesting. It's a documented DNS encryption scheme. What's this about?

Steve: Well, so this is only for OpenDNS users.

TOM: So you have to be using the OpenDNS service to serve your domain names?

Steve: Yes. So what they've done is, this is not an Internet standard. This is not DNSSEC. And DNSSEC, which is still not widely available, we've covered it in detail in previous podcasts. It provides spoofing detection and authentication, essentially secure signing of DNS records, which allow somebody who makes a DNS query - now again, this is DNSSEC. This is not what I'm talking about with OpenDNS and DNSCrypt. I wanted to draw the distinction, though.

TOM: Yeah, we just want to understand what DNSSEC is so we could tell the difference; right?

Steve: Right. And so it provides a means of signing DNS records so that the recipient of a DNS query can absolutely, with cryptographic security, verify that the record they received is the record that was signed by the record's owner, the authoritative owner and distributor of that DNS record.

TOM: And not being served by Immigration & Customs Enforcement, say.

Steve: Precisely. Not altered in any fashion. So what the OpenDNS company has done is - and this is just sort of random. I mean, I'm not sure why they did it. Certainly they're looking for ways of inducing more people to use their service. And so here's a new feature. The way to think of this is...

TOM: I think that we don't - and we should point out, we don't have DNSSEC rolled out yet. Right?

Steve: Right. And it's, I mean, the root servers were only just recently signed. So those records are signed. But in general, the availability of DNSSEC is virtually nonexistent. It's like IPv6.

TOM: So maybe this is a Band-Aid to get us by until DNSSEC is fully rolled out? I don't know.

Steve: No. This is...

TOM: Too farfetched?

Steve: The way to - yeah. The way to think of this is a feature that one DNS provider, OpenDNS, has now in beta only for Mac users. So if you're a Mac user, and you are using OpenDNS, or you like - or you're a Mac user, and you like the idea enough that this would move you to OpenDNS, then you've done what they really want, which is moved to OpenDNS. So what this is, is this is point-to-point encryption of DNS records between a Mac running the OpenDNS DNS client, meaning you have to run something special on your Mac side in order to make this point-to-point encrypted connection to the OpenDNS servers.

TOM: So this isn't going to work on a router.

Steve: Well, correct. It will go...

TOM: You have to run it on the actual machine.

Steve: Yes. So the idea is, normally your Mac would - and that's a very good point. You could not have your router configured to use OpenDNS and then have your Mac receive - have your router do the queries on behalf of the network. Now, most routers are not generating DNS queries themselves. They're merely offering the DNS IPs to the machines on the network, and then those machines make the queries. So that would still work. So the idea would be, if you've got a Mac, and you're running this OpenDNS beta for Mac users only, then a low-weight, that is, a relatively inexpensive-to-establish encrypted query is made to OpenDNS, and then an encrypted reply is received.

So what does this do? This, more than anything, this gives you privacy, which is something that even DNSSEC won't do. DNSSEC gives us verifiability and authentication, but specifically not privacy. DNSSEC does not encrypt. Even when it's eventually deployed, it's not encrypted. So this does give you point-to-point privacy. This means that in an open WiFi environment, like in a Starbucks, where you don't have any encryption key on your wireless, anybody sniffing your connections will be able to see all of your DNS queries. They'll see what sites you're looking up by your query and the IPs of those sites coming back to you.

TOM: They just can't interfere in that case, right, and perform a man-in-the-middle attack because it's a signed domain name.

Steve: That's one of the things...

TOM: Probably.

Steve: ...this is supposed to prevent. And actually I have a question about whether it does, that it prevents man in the middle, because I don't see how it can. I haven't looked at it in depth because it's in beta at the moment. And, for example, you have to reenable this every time you boot your Mac. They don't even have it set up so that it remembers to stay on.

TOM: Oh, well, that's annoying.

Steve: That's one of the things in their FAQ. They said, quote: "If you have...." Oh. They said the beta must be reenabled after every reboot. So anyway, so I wanted to let people know this is there. It's pre-release. I wanted to make sure people knew that they had to turn it on again every time they use it, that is, every time they boot their Mac. But for people who are interested in privacy for DNS, this is available. So it's like, okay. It is what it is, an interesting service that OpenDNS offers. I guess if you're a Mac user, you're already using OpenDNS, you might give this a try and see what you think.

TOM: Do you have any idea why it's only for Mac? It seems like an odd way to go.

Steve: Yeah. Maybe they're Mac people, and they just haven't gotten around to doing it for Windows. It might be that the Mac platform, being a little more easy to maneuver - for example, getting this into Windows might be more difficult. I just don't know.

TOM: So let's move on to the new Zeus banking trojan, which has got a nice little one-two punch, apparently. This is not your father's trojan.

Steve: Well, Zeus has been a problem for quite a while. It's become the premier ACH banking trojan, meaning that people download it into their machines by mistake. It lives in their system, watches them doing banking transactions, and then hijacks their banking credentials. So that's actually the way this thing runs.

What the FBI has reported they are now seeing is that the latest variants of the Zeus trojan will perform a fraudulent transfer in order to obtain moneys from one or more individuals, and then - so that's the "one" part of the one-two punch. The second part is they will then launch a distributed denial of service attack on the banking website in order to pull the site down. The reason they do that is that this gives them a larger window of opportunity in order to get the funds moved to somewhere where they cannot be reversed.

So the way this normally fails, that is, the bad guy's ability to transfer money completely out of a mode where it cannot be reversed, is that the end user may notice something going on. I mean, they're in the process of doing their banking with the website. So if the trojan gets the credentials and moves the money away, if the user, in doing their banking, checks their balances and says, wait a minute, it looks like I don't have as much money in my account as I thought, well, it's just because it just at that moment was transferred away. So they then contact the bank and complain that there seems to be something wrong. The bank pursues this and sees that there was a fraudulent transfer and reverses the funds immediately.

Well, so what now we're seeing is that immediately after the fraudulent transfer is made, a denial of service attack is launched on the bank to prevent the valid users from being able to get to the website to see that their money is no longer there. So it steals their money and then crashes the site so that they can't tell that their money has been stolen.

TOM: Now, you could theoretically still call in, phone banking, and check your balance. But I guess the idea is people don't do that as often as they might check online. And that delay is all the bad guys need?

Steve: Well, yes. But also it's one thing for the site to affirmatively tell you that you have a zero balance, and it's another thing for you just not to be able to get to the site. Because, I mean, the Internet's still a little flaky enough that if someone tries to use the website, and it doesn't respond, they think, oh, well, I'll come back later and try it again.

TOM: They're not usually going to pick up the phone and say, oh, something's wrong, I'd better call in and check my balance. They're just going to assume that their Internet connection is down.

Steve: Exactly. Like there's something wrong with the banking website, and they'll check it again later. Now, okay. The other thing I've picked up on in doing a little bit more background on this was an interesting new trick that the bad guys are using because they have the problem of transferring the money somewhere, and then what do they do with it? How do they get it out of reach of the banking system so that it can't be grabbed back?

One of the tricks now that's being used is that the bad guys will contact high-end luxury jewelry stores and commit to purchase precious stones or, for example, luxury watches that are very expensive. And they will say - they'll arrange the purchase, and they'll say, okay, we're going to wire the money to you, and then we'll have one of our people come and pick up the merchandise. And so the jewelry stores says, yeah, that's fine. As long as you've wired the funds, that works. They will then transfer the money from one or more individuals who have been infected with the Zeus trojan into a composite account, and then from there to the jewelry store.

They will then verify with the jewelry store that the funds are available. And then one of their so-called "mules" will come in and pick up the physical merchandise which has now been paid for from this composite money transfer, wiring the funds into the bank, that then - everything's been paid for. This guy then comes in, the so-called mule, and absconds with the merchandise. Then, when the law enforcement gets involved, they realize what's gone on and will reverse the transaction, but the jewelry store has already allowed the goods to leave its control.

TOM: So that DDoS is there to give them time to get to the jewelry store and make off with the jewels.

Steve: Exactly.

TOM: Wow. This is like a 1950s heist film.

Steve: It's just amazing. And this is what's actually going on now every day.

TOM: Yeah. I expect to see, like, Frank Sinatra or George Clooney in the movie version of the Zeus trojan.

Steve: Right.

TOM: Somebody should option that.

Steve: So there's been, as you - I'm sure you've been covering this in TNT. There's been lots of discussion about this Carrier IQ.

TOM: Oh, yeah. I've been looking forward to getting your perspective on this because we have been covering this quite a bit. For anybody who doesn't know, Carrier IQ, a program that a researcher discovered was existing on the phones - sometimes it's an app, sometimes built into the operating system. And on this particular researcher's HTC phone, this guy named Eckhart, he found that it was - it seemed to be monitoring your keystrokes, the contents of text messages. Turns out that it's probably not as nefarious as it looked at first. Rebecca Bace was hired by Carrier IQ to look at it independently. She said it's not doing any keystroke monitoring. Dan Rosenberg independently looked at it, and he's had a good report about what it does and what it doesn't do. So what do we need to worry about here, Steve?

Steve: Well, so - okay. This reminded me, the whole episode reminded me of exactly what originally happened with the first discovery of spyware. And Leo likes to remind people that I'm the person who coined the term "spyware."

TOM: Catchy, by the way. Good job on that.

Steve: [Laughing] I remember it was something that I learned about had been installed in my machine without my knowledge or permission. And this is, I mean, this is, like, old school. This is a long time ago. I learned about it because I was beta testing ZoneAlarm, which at the time was the very - it was the only personal firewall that did outbound monitoring.

TOM: I used ZoneAlarm for years. It was a great firewall. And free.

Steve: Yup. And it was, like, two days after installing it, I got a popup saying that - and the details of the - wait a minute. Aureate, it said Aureate, A-u-r-e-a-t-e, was the software that had been installed in my machine. Well, this was not malware, technically. This was so-called "advertising" software which had been brought in by WinZip or Zip for Windows or one of - it was one of the Zip utilities. And I had installed that. This thing had been installed without my knowledge or permission. And it was then monitoring what I did and sending this information back. And the ZoneAlarm firewall caught it doing that.

So immediately, it's like, wait a minute. What do I have - what is this? And this, again, this was a very, very early piece of software. So I made this news public, and it brought a firestorm of reaction down on this Aureate company and on the people who were carrying this into people's machines. Well, so this is very similar to the same thing that has happened because the question is, is it the handset maker's fault for this being there? Is it the carrier's fault, like T-Mobile or AT&T? Well, and even Apple has this installed on their iOS devices.

TOM: Well, there's reference to it in iOS 5, and they say they're going to get rid of it. They don't use it anymore, and they're just going to get rid of all trace soon.

Steve: Right. In an update to iOS it'll be completely removed. So executives at Carrier IQ have said that their monitoring software gathers information about web usage, as well as when, where, and to what numbers calls are made and text messages are sent, but not the content of text messages. I did see something that indicated they are capturing URLs that smartphone web surfers are surfing to. So that's some concern. But they're also saying that, you know, they care about text messages only inasmuch as did the message go through? Did the message not go through? If not, then what cell tower is the user using, that sort of lower level, service level stuff. So multiple class action lawsuits have been filed against Carrier IQ, just as multiple lawsuits were filed against this Aureate company. People mostly are upset with - they're concerned that they feel no one told them this was going on.

TOM: Right. It was...

Steve: It's been installed behind their back, without their knowledge and permission. And suddenly it's like, wait a minute, I'm being spied on by my phone.

TOM: And Carrier IQ and the carriers themselves are caught with their pants down here because they've been doing this for decades without anybody knowing or caring. And I tried to explain this on TNT yesterday. You think about your old handsets from the '90s; right? They were doing all of this stuff. They were logging all of this information because the handsets were sort of considered part of the network. They weren't yours. They were a piece of equipment you used to access the network. All this information, even URLs, but things like when a text message is sent, what cell tower, all of that stuff is stuff that has to be part of the network. The carrier has to know it to complete your calls. They know all the phone numbers you dial, too, but they have to. And so what Carrier IQ was doing is collating all of that information and providing some diagnostic tools. And so back then, with features phones being the only phones, nobody thought twice about it.

But now, smartphones are ours. They do a lot more. We entrust a lot more information to a smartphone - banking information, passwords, private messages that we didn't really rely on phones for in the past. But this software is still there in the background, and it's still doing what it was always doing is trying to collate diagnostic information. And it sounds, from what I've read, that it's not doing really, really nefarious stuff like keylogging. But it is passing along a lot of information about a device that is very personal now without your knowledge. And I think that's really the big thing is that they don't make it clear what it is that it's running, and they don't give you a chance to opt out or opt in, either one.

Steve: Right. It's funny because what you just said reminds me of the model and the example that I've often used with personal computers. It used to be the case that people would say, oh, I don't worry about my PC security because I don't do anything important on my PC. I just use it for surfing the web and doing email, but that's all. And my response had always been, well, today that may be true. But the world is going to be pushing you towards more use, more comprehensive use of that same device.

And what will happen is that your bank will start encouraging you to do your banking online. And so there'll be some creepage of your use in a direction where security really does matter. And if you never bothered to give your machine a good password, then you slowly start using it for more critical things, then you're beginning to increase your exposure over time. And so it's really necessary, if you say deliberately, "I don't care about the security of this device," to be sure that you never do anything with it where you do care about the security, even in that instance.

TOM: Yeah. I think that Carrier IQ probably isn't the bad guy here. I think they're just behind the times. They did a really bad move by trying to use a cease-and-desist copyright infringement notice...

Steve: Oh, yes.

TOM: ...to shut down Eckhart when he first put out his response. They've since scrambled to try to do the right thing, but it's too late now. But the issue really revolves on should they change with the times and should they make it clear. Sprint actually has, in their terms of service, that they can collect this kind of information. They just don't disclose how they're going to collect it, that it's being collected all the time. It's not satisfactory to me that it's disclosed as well as it should be.

So I think Carrier IQ is really trying to figure out how to do the right thing at this point. But they're going to have to change the way their software works. To satisfy me, they're going to have to say, look, this is going to be removable. We're going to have to alert people that we're collecting this information up top, and what information we're collecting, and give you a chance to opt in and say, yeah, I would like the network to work better, and I would like to contribute to that. As long as I know that the contents of my text messages, the contents of my emails are not being collected by somebody, I personally am okay with them collecting all the information they already know because they have to know it to make the network work, and using that to do better diagnostics.

Steve: Well, what I think will happen is that all smartphone users will get an education about what their carriers are collecting. I think that that's been part of the problem. As you say, this has been going on in one form or another forever. And as our platforms evolved, they were able to enhance and increase the sophistication of software running out at that endpoint.

And so certainly everyone would agree that the contents of our email and the contents of our text messages should not be spied on or spy-able by some random third party. As you pointed out, Tom, this stuff all goes through the carrier. I mean, the carrier knows about the content of our text messages because text messages are not encrypted, and they're a service of the underlying platform. So it was this idea that some third party that we weren't aware of was involved.

And so I think that over time it will just become understood that there may be some subcontracting going on that the carriers use for collecting this data; that it comes with the territory; and it's about cell towers and service-level things, not content-level things; and that people - it'll just be - we still may not read the fine print. But there'll be, as a result of all this kerfuffle, sort of like, oh, yeah, I heard about them, and it ended up not being a big deal. And that's just what happens with smartphones.

TOM: But let me be clear. I want the ability to turn this thing off at a moment's notice if I decide, you know what, Carrier IQ, I heard that they're starting to collect some things that I don't like. I want to have control over it as the user. I want to be able to opt out of it. And right now the only way to opt out of it is to install a new ROM on your phone, and that's not acceptable.

Steve: Yeah. I don't think opting out will be available. I think your carrier will say, if you want service from AT&T, this is part of the service. I mean, it's built into the phone. It's what they require in order to offer their service. I doubt that turning it off will be an option.

TOM: Well, you may be right about that. They can't - in some cases they do it as an app instead of built into the operating system. So they have the capability of doing it. Whether they will or not, that's a whole different question.

Steve: Right.

TOM: Let's move on to we've got some miscellanea. Jakob Nielsen has weighed in on the Kindle Fire and other seven-inch tablets. What does the master of web usability have to say?

Steve: Yeah, well, first of all, for people who don't know his name or who recognize his name, and you obviously do, Tom, this guy, J-a-k-o-b N-i-e-l-s-e-n, has long - he's written, like, some of the classic texts on website usability and accessibility. He's regarded as a major UI guru. I, in my Twitter stream recently, @SGgrc, I tweeted the link to a very nice review that he did, sort of a usability study, only a small number of people, but he felt like he got some very good results from it. So you can just go Twitter.com/SGgrc, and you can easily find the shortened link to his review where I refer to this.

And I just wanted to share with our users the upshot of him looking at this was to conclude that seven-inch tablets were too small for convenient use of non-mobile websites. So websites that have been designed for mobile screen sizes were fine. But just wandering around the web in general, a seven-inch screen was unusable for non-mobile websites; whereas the iPad or a larger screen, a 10.something-inch screen, was large enough for non-mobile websites. So I just sort of thought that was an interesting conclusion for them to come from, that I wanted to share.

TOM: Yeah, the Kindle Fire, I found, is really good for reading books. I feel like it's light, you know, it's lighter than an iPad. But it's got a nice resolution, and it's kind of paperback size. But I have not found it to be great for anything else. Not for me yet. And maybe - I haven't been on a plane with it yet. So we'll see if I like it for video. I have a feeling I want that bigger screen. I want that 10-inch screen when I'm watching video.

Steve: Yeah.

TOM: Also on your Twitter stream - people should follow @SGgrc - a great review of current eInk readers.

Steve: Yeah. There was a very nice review from the guy who developed Instapaper. So he's well known and been around for a while. What I liked was that he took multiple brands of eInk readers and really did a very nice side-to-side comparison. So anybody who's been wondering among the various eInk readers, that is to say, the reflective display, not the LCD tablets...

TOM: Right, we're not talking about the Fire or the NOOK Color or any of those here.

Steve: Precisely. Those readers that are eInk and monochrome, he looked across them all. And I'll just give you his - his result was that he liked, of all of them, the smaller Amazon Kindle the best. That is what is now just being called the Kindle. Not the Touch, and not the so-called Kindle Keyboard, but just the little Kindle.

TOM: The one that's $79 if you buy the ad-subsidized one. It's super cheap.

Steve: Precisely. And my feeling is that it's actually a little small. It's sort of tough to hold it because there's no, like, paddle area at the bottom, which I enjoy on my Kindle 3, or the so-called Kindle Keyboard. I sort of like that better. I mean, it's nice because you can really put it in your pocket and carry it around, so it is small. I almost think they undershot on the size end.

And speaking of Twitter followers, yesterday I crossed the 25,000 followers mark.

TOM: Congratulations. That's fantastic.

Steve: So for all of the podcast listeners who are following me on Twitter, I wanted to just send a thank you. I tweeted a thank you this morning, telling everyone that I will continue posting important and useful tidbits whenever I'm able to.

TOM: Well, here, yeah, we've got two examples here: the Jakob Nielsen stuff, the reviews on eInk readers that you could have found out there. See, I follow. I'm one of the 25,000, Steve. I follow. It's great stuff in that feed.

Steve: Well, and I did have a nice note from a Paul Smith, who - actually, this is a different kind of testimonial than I normally post. He said, "SpinRite Saves Movie Night." And he's in Ipswich in the U.K. And he said, "Hi, Steve and the support team. Sunday night is movie night for us. However, the disk drive in our NAS, our network attached storage box, started going AWOL. So I took it out and connected it to a PC and ran my copy of SpinRite - first use since I ordered and purchased it back in 2010. Anyway, at about 16 percent I got a divide overflow error and a red screen.

"This being Sunday, I thought that was it. But no, an email to GRC Support was returned in minutes, and changing a setting in the BIOS as recommended by Greg completely cured the problem. Within the hour, SpinRite had completed its magic, and movie night was back on. Of course, the drive then worked perfectly. Thanks, Steve, for the great software. And you can pass my thanks to Greg McIntyre from your support team. And also thanks for the Security Now! show." So I just wanted to let people know that they're getting something for their money when they purchase our software. We really are here, maybe not 24/7, but at least 7.

TOM: Yeah, yeah. Well, you saved movie night. That's important. That's the kind of stuff people feel much better about, when it's sort of like, gosh, this is something I was really looking forward to, something I really enjoy. I didn't think I was going to get a response. Good customer service, that's just as important as the good product.

Steve: Exactly. That is what you're buying.

TOM: All right. We've got a dozen questions today for Listener Feedback #132. You ready to dive in, Steve?

Steve: You betcha.

TOM: Let's go with #1 from @Jason_JW in Nacogdoches, Texas. Used to have a union organizer friend of mine lived in Nacogdoches. He says: Steve, any recommendations for PC performance diagnostic software? My boss complains the PC I recently built is slowly down. Thanks.

Steve: Okay. So this was something that was tweeted, interestingly enough. I wanted to mention that I do kind of keep an eye on my Twitter stream. So if someone wants to mention @SGgrc, I will normally see that. I get a lot of great tips from our listeners, and so that's really a great source of information which I very much appreciate.

So I thought about Jason's question, saying that a recently built machine was slowing down. And he was asking me, like as you read, recommendations for PC performance diagnostic software. And I thought, okay. I'm aware of a number of different things that have been around. For example, there was something I ran maybe six months ago that did an audit of my machine's startup, so that I was able to see what was going on as the system was starting. But I sort of pushed myself back a little bit. And I said, you know, one thing that almost anything you install does is make a mention of itself right in the control panel under Add/Remove Programs.

And so what I normally do when I first encounter a system which is running sluggishly, is rather than installing any other kind of third-party stuff, I just jump over to the Control Panel and look at Add/Remove Programs. And it's very revealing to see how much crap, frankly, a machine can acquire, in someone else's hands, in a couple of months. There are some people who just seem unable to resist installing things. And very often these are things that are installed once.

So many programs now that you install want some time from your system when they're starting up. They want to run in the background. They want to run a chunk of themselves in the background. I'll see that, like, for example, things I rarely use, like for example Microsoft Office Suite, will have a quick startup thing that it's running whenever it starts up. And it's like, okay, do I really need all of that? So I would recommend, before you do anything else, just go to the Control Panel, look at Add/Remove Programs, and scroll through, maybe sit down with the person whose machine this is and say, okay, when was the last time you used that? Do you really need that? And what about that one? Very often these are things they installed once and then forgot about, haven't even used since. And just getting rid of all that junk can really make a difference. So I think that's what I would recommend is just use - just go through and look at, sort of challenge them for the things that are installed, do they really need those things.

TOM: Yeah. I think I can add a little bit to that, which is you don't really need software to do diagnostics so much as a good checklist of inspections. And that's what you're talking about. And look at MSConfig, too, see what's starting up because that can slow stuff down. And I just reinstalled WinDirStat recently, for the first time in a long time, which goes through and shows you graphically what's taking up space on your hard drive, which is another good way to find out - because when you see that list, they show you the size of the program. But seeing that graphic layout, it really says, oh, wow, I've got a lot of stuff in that directory. Maybe if I just clean that out and then do a defrag, that helps. Good stuff.

Question #2 comes from CyberAdminDude in Montreal. He wants to know how to stress the importance of security to "n00bs." He says: I am a catchall server/programmer guy for a small company that runs a website. I was wondering what the easiest way to stress the importance of good practices to people less in the know is. When I talk about them, I get blank stares or NyQuil-like reactions to my rants. The person who is my commander isn't really tech savvy, either, but knows the buzzwords, LOL. So I don't have the influence/power to make changes myself. We have passwords like "apples" for our phpadmin page and FTP passwords like "Workus3r." Is there a website with a collection of testimonials or horror stories I can send to my coworkers to help them understand the threats?

Steve: So, okay. I like this question because it's one that I see a lot in various forms. We have our listeners, who understand what's going on with security, understand that the fundamental importance, that is, that security plays in today's world. And so many people are frustrated with people who just don't give it a thought. And they say, exactly as CyberAdminDude does, how do I get this across? How do I explain that this is really important?

And the best advice that I have for people is try to connect the notion of cybersecurity to aspects of security that people do understand, and that's real-world security, physical security. Would you deliberately leave your front door unlocked? Would you deliberately walk away from your car with the windows rolled down while there's stuff in your car that would be valuable to somebody else, that might be available for stealing? I think that the problem that people have with cybersecurity, not understanding what's really going on, is that it's too virtual. They just don't think of it in the same way they think of physical security. We've talked about how no security is perfect. Yes, locking your front door doesn't prevent people from breaking your windows. But locking your front door is better than not locking your front door. And closing your windows is better than leaving them wide open.

So I've sort of thought about this a lot, and I think that the best advice is to try to relate this to security that people do understand. And everybody understands physical security, real-world security, security that you grow up gradually acquiring an awareness and an appreciation and an understanding of the importance of, over time. And try to explain that the same issues in the real world do affect the virtual world. And so it's important to take these precautions. Just leaving your windows down doesn't mean someone's going to steal from you, it just means they can much more easily. Similarly, using a password like "apples" doesn't mean somebody is going to crack into your phpadmin pages, but it means they can much more easily. It's exactly analogous.

TOM: In fact, if you want to look for horror stories, I would say just look at the news. Pretty much every week you'll find two or three. And I would suggest using those in conjunction with what Steve's saying, which is like, would you leave your windows open? And if someone did break in and steal your car because you left your windows open, how would you feel? Do you want to be that guy who left the password as "apples," and then a hacker breaks in and it makes national news, like this water pump that allegedly got hacked? Do you want to be that guy? Try to relate that sort of like you will be, and what's your defense going to be? Oh, I tried to secure anything? No.

Steve: Right.

TOM: All right, Question #3. John Palmer in Washington, D.C. has a question about Carrier IQ. A lot of people have a lot of questions about Carrier IQ. On December 1st he tweeted Carrier IQ...

Steve: No, wait, on December 1st "I" tweeted.

TOM: Oh. He says: Steve, on December 1st you tweeted "Carrier IQ not a rootkit. It's commonly installed carrier feedback firmware for monitoring handsets. Comes with the territory." But, he says, it clearly is a rootkit, isn't it? So is it a rootkit or it's not a rootkit? Why isn't it a rootkit?

Steve: Well, okay. So it was interesting how much controversy this stirred up because I was seeing people using the term "rootkit." Actually, people were sending me in my Twitter feed all of this commentary about Carrier IQ being a rootkit. And so I tweeted, "It's not a rootkit." And that generated still more controversy. So I thought I would just take this moment to say to our listeners, to remind people, a rootkit is not software that you don't like. I mean, okay, a rootkit is software that you don't like, but it's not defined as software that you don't like. A rootkit is something that a malicious third party installs in your machine through a vulnerability, and it's not something that your OS vendor approves of, not something that you approve of. It was maliciously installed.

So we now understand that, even though people may not like the idea that there is this monitoring technology which was offered by a third party, it's not a rootkit because it was there when you got the phone. I mean, it is - somewhere in the license agreement it says it's going to be there. And it may have been hidden from you deliberately. In some cases that's so that it's not mixed in with the applications that you see and you're able to run and install and remove and stop running and so forth. But it's also - it's part of the firmware. So I just wanted to draw that distinction. Rootkit, again, rootkit isn't defined as stuff you wish wasn't there. Rootkit is malware which was installed and hides itself. So, yes, Carrier IQ is preinstalled and hiding itself, but that's not a rootkit.

TOM: And does a rootkit have to have some sort of kernel association?

Steve: It generally does, although that's normally, I mean, "root" means, in UNIX parlance, god-level rights, full rights to the system. And it hides itself by using kernel hooks in order to prevent itself from being found. So this is just sort of - the Carrier IQ isn't a rootkit in that sense. I mean, it is down in the OS, in the kernel, monitoring what we're doing, but with full knowledge of the provider of the handset. And that's the difference.

TOM: Sony rootkit was not - Microsoft was not aware that the Sony rootkit was inhabiting Windows.

Steve: Right.

TOM: Web791 puts it well. He says this is more like a root canal, just something painful that you don't want there.

Steve: [Laughing] That's pretty good. I like that.

TOM: But the dentist knows that he's doing it to you. All right, let's...

Steve: Well, and you're also paying him for it.

TOM: Yeah, true. You're paying the carrier for your phone, too.

Steve: Right.

TOM: Let's get to Question #4. Aaron in Bend, Oregon is not alone. He's baffled by bandwidth. So many people are. He says: Hi, Steve. I'm confused by bandwidth. Everywhere you look, the only measure you find is Mb/s. Is it really that simple? Does that mean that my office of 30 people with a 1.5 Mbps connection is truly the same as my home DSL connection of 1.5 Mbps? Shouldn't there also be a capacity measurement? I'd like to think that the office connection is like a multilane freeway with a speed limit of 65, and my home is like a rural two-lane road with a speed limit of 65. Both have the same speed, but different capacities. Could you clarify?

Steve: Well, it's an interesting analogy which is a little bit confused. But we'll have some fun here with some visuals. We recently on the podcast talked about the way bandwidth is throttled, that is, I think it was a Q&A maybe either two weeks or four weeks ago. Somebody said, how is it that bandwidth is differing for different users, if all the packets on the Internet are moving between routers at the same speed? Which is the case. The routers out on the main Internet backbone are treating everybody the same. It is the so-called "last mile," which is the term used for the ISP to you connection, where the ISP is hooked into the main backbone of the Internet, pulling, able to transact traffic with all of the other major providers on the Internet. But then your ISP is sort of your portal, your connection out onto that super high-capacity highway. So what your ISP does is limit sort of the average speed that you're able to move traffic to and from the Internet.

Now, the reason that this wide-lane, multilane freeway analogy, where all the cars are moving at 65 miles an hour, versus a rural two-lane road with the same speed limit, the reason that's sort of confusing is that the way to think of bandwidth would be in cars per second, for example. So if all the cars are moving at 65 miles an hour, but you've only got one lane in each direction on a rural road, then that's going to limit the number of cars per second that are able to travel. But if you had a multilane freeway, where each lane was moving at 65 miles an hour, because now you've got parallel lanes, the number of cars per second is much higher on a freeway, which actually is why freeways work.

TOM: So in other words, the miles per hour in this analogy is equivalent to the speed of light.

Steve: Yes.

TOM: Because all your bits are moving at the - well, and maybe not on copper. But essentially the bits are moving the same speed all the time.

Steve: Right. Right. And so we have the notion of packet rate, where an ISP limits you to a certain packet rate, that is, the number of packets per second you're able to move. The packets themselves may move very quickly. In fact, the actual packet, the actual rate at which the bits move would be the same for a low-bandwidth cable modem user and a high-bandwidth cable modem user. It's just that the high-bandwidth cable modem user gets more packets per second, where the actual packet rate is the same.

And now, finally, stepping back and looking exactly at what Aaron was talking about with his office connection at 1.5Mb and his home connection, what really happens is, if you were to monitor a single user's 1.5Mb connection, it would be very burst-y. That is, he clicks a link, and that sends the link off. And then the page comes back. And then he sits there and looks at the page, scrolls a bit, clicks a link. Again, that happens. So he may have a 1.5Mb connection, but he's not using it, that is, his utilization of it, because he's just one person, is probably relatively low. But if you put his 30-person office behind that same 1Mb connection, you're going to see it is saturated by those 30 people.

So now you've got 30 people, all clicking links and probably waiting a little bit longer, that is, their traffic is, like, lined up one behind each other, really packing that 1.5Mb connection much more densely than Aaron sitting by himself at home. So even though both situations, a 30-person office and one-person home, may have a 1.5Mb connection - and this is what he talked about when he's talking about capacity. It's actually the utilization level of that link at that speed is probably much higher when you've got 30 people all fighting each other for access to the Internet over that same relatively low-bandwidth channel, compared to 30 people. Essentially, if they were all trying to use it all the time, they would be 1.5Mb divided by 30 would be their shared level of access.

The good news is it can function pretty well because most users who are not just downloading monster files, but interactive users, are inherently burst-y in their access. They click a link, get a page, look at it, click a link, get a page, look at it. And that's the kind of access that 30 people could share just by sort of interleaving their access.

TOM: Well, and this is why bandwidth caps are not bandwidth caps. They're data usage caps.

Steve: Right.

TOM: And there's been some good studies recently that show that they don't do anything to alleviate congestion. And it fits right into your analogy. A high data usage person, someone that's going to run up against that 250GB cap from Comcast, for instance, is just someone who drives a long way on the road. He spends a lot of miles on the road; right?

Steve: Right.

TOM: But he doesn't contribute in any relevant way to the congestion. The congestion happens when you have a lot of cars trying to use that highway, and you reach the capacity. And I've got the names here: Herman Wagter and Benoit Felton talked an unnamed DSL company into allowing them to look at their customer data in five-minute increments, and it bore this out. The so-called "bandwidth hogs" who use a lot of data did not contribute in any relevant way to the actual congestion that the DSL provider was experiencing. And it's because of that burst-y way that things happen that you're talking about. It's just it's numbers. It's sheer numbers of people who use the service.

Steve: Yeah, it would be like, using your analogy, it would be like if a huge number of people wanted to only drive one exit on the freeway. Well, they still would have a problem, even if they're only going one exit. It's not how far they go, it's how many of them are trying to go per unit of time.

TOM: On that same stretch of road, exactly. All right. Question #5 from Kevin Odell in Levittown, Pennsylvania, wants to know how he can test router guest network security. First he wants to thank you for Security Now!. He really looks forward to every week and has never missed an episode. He says: How can you test guest networking is properly segmenting the networks? I know you mentioned the Airport Extreme does it properly. I recently purchased a D-Link DIR-655 router and have had a lot of problems using AirPlay between my iPhone and my Apple TV 2. Recently a friend was at my house and said, "Cool, I can see your Apple TV." I knew immediately there was something wrong with the router since he was on the guest network, and the Apple TV should not be visible. I did a firmware upgrade, and still the same thing. I turned off the guest networking, and AirPlay works perfectly. I went through three levels of tech support, and they finally passed it on to the engineers at D-Link, and they still can't explain it. I just wanted to make your listeners aware, don't trust guest networking unless you know for a fact that it's working properly.

Steve: So that's a great piece of wisdom. One of the things that I hear frequently as the person who offers ShieldsUP, which has been used for untold number of years by people to check their security, what I'm hearing all the time is they'll believe they're secure, that they'll believe their ports are closed or that they know the way their system is configured, they'll check it with ShieldsUP, and ShieldsUP will find a port which is open that they were unaware of. And so it's really the case that testing is the only way to know for sure what's going on.

Now, I did a little bit of research on D-Link in particular, since that's what Kevin was referring to. And there is a setting on the guest network configuration, which you have to explicitly disable, which allows routing between the guest network and the main home network. So you absolutely want to make sure that's turned off. If not, then the router will allow traffic to cross between guest and main network. And even so, you want to test. The way to test is as simple as just trying to ping the router. I mean, I'm sorry, try to ping the machine on either side that are in these networks which are supposed to be isolated.

Take a look at the IP, for example, that a machine has been assigned over on your home network. And then, from a machine on the guest network, just use the ping command. Open up a command window and type p-i-n-g, space, and the IP address of the other machine. See whether you get a response. You hope not to, if that machine is not accessible, and you want your networks to be isolated. And I would just say, if you're using a guest network, regardless of what router you have, look carefully at those settings and make sure that the router has been instructed to isolate those networks. Because it makes sense that there would be the option because in some cases you might want your guest, for example, to have access to your Apple TV device. In other cases, you definitely don't.

TOM: Yup. Question #6, Luis in Spain has an interesting observation about packet TTLs, that's Time To Live. He says: Hello, Steve. About the TTL or hop count, I want just to inform that, when a packet goes through a MPLS VPN for a customer, the provider doesn't touch the TTL. So a lot of times there is not a TTL change. Most packets have more hops than we can see on the traceroute. Just to inform you of that. Kind regards, Luis.

Steve: I loved that. It's something that we had never talked about before in our - you were talking about the How the Internet Works series.

TOM: Yeah, yeah, love that series.

Steve: And one of the things we covered carefully was this notion of the TTL, the so-called Time To Live, which as Luis mentioned has been renamed "hop count" in the IPv6 spec. They wanted to make it clear that it wasn't a time measured in seconds or any temporal sense. It was actually decremented per router hop, and so simply called "hop count."

But the point he makes, and it's a really good one, is when you are in a virtual private network tunnel, then your traffic, as it moves from router to router, the external tunnel packets will have their TTLs decremented. But the packets moving through the tunnel don't see router hops at all. So there will be no TTL decrementation for tunneled packets, that is, packets that are being carried by the VPN tunnel, which is something I had never mentioned before. And I thought that was a neat observation. So if you were, for example, to do a traceroute from a point to another point, not through a VPN, you would see every hop count shown by traceroute that the packets made.

By comparison, if you did the same traceroute, but some portion of the transit was carried by VPN, then you would see none of the TTL decrementation or the IPs of the routers which a traceroute would normally show you, thus tracing your route, until your packet emerged from the other end of the tunnel, then made any additional hops it needed to, to get to its destination. So you are in fact blinded by the tunnel. They don't get seen. I thought that was just a very cool observation.

TOM: Yes. That means it's working, is what that means.

Steve: Yes, exactly.

TOM: Christopher S. Bates in Central Valley, California asks about special characters in passwords. He says: I've been catching up on your podcasts here in the last few weeks and remember you mentioning a problem with some sites disallowing special characters in usernames and passwords. That's one of my bugaboos, too. I completely agree with you that this is horrible practice. I have noticed, though, and tried to get changed, that my bank follows this practice for logging into their customer web portal. I have reported this as being an issue many times to their online suggestion box/email system, but it has never been corrected.

I understand that this is probably something that I shouldn't worry about, considering they do enforce other security measures. But it is something I feel should be enforced on all sites and systems, especially those dealing with financial institutions. Do you have any thoughts on this situation, or any suggestions on how I may get them to change their policy?

Steve: And you want to keep reading.

TOM: Well, okay. I wasn't sure if we were - he says: I left this in case you don't want to read it over the air. I'm talking about Chase banking. I just wanted to make it clear to you this is a large institution, not a small credit union.

Steve: Yup. I thought we ought to put Chase's feet to the fire.

TOM: All right, Chase, you've heard it here.

Steve: Yeah. Okay. So we've talked about this before. And since we last did, I verified some of the rumors I had heard. The reason this seems to afflict banks, annoyingly, is that the way the web evolved was that, unfortunately, web-based front ends were put in front of existing old-school mainframe banking back ends.

TOM: Aha.

Steve: And so, I mean, it is slipshod, and it is sloppy, and there's no excuse for it. But there is an explanation for it. So I don't mean to be excusing this behavior at all, merely understanding and explaining it. And so it is because once upon a time the login technology for mainframes wasn't very secure, and it only allowed alphanumeric passwords. And so what happened was that exactly that technology was just sort of pushed out onto the Internet so that users are logging in, in the same way over the Internet that they once logged in directly at a mainframe terminal.

Now, there's nothing that would have prevented a much more sophisticated and secure front end to provide essentially separate web accounts which would then have an identity to the mainframe, so that users could log in with all kinds of extra security, multifactor authentication technology that the back end didn't ever need or think to support. And then, if all that succeeded, that would then log them in using old-school and private alpha-only login, so that they had to authenticate with a much more secure front end.

That's not what happened. And that's why we keep seeing banking institutions having among the worst web-facing login security of any. It's because of the legacy of mainframe login that just got surfaced out onto the web page. So the bad news is I don't think, no matter how many times Christopher and any of our other listeners complain, we're going to see any change. It will end up being a legislative requirement imposed by law. At some point they will say minimum password length and large character set must be supported, and they must be case sensitive, and so forth. It will be that kind of legislation which finally enforces banking institutions to say, well, we're going to have to spend some more money. They just don't want to spend the money.

TOM: Yeah, they don't want...

Steve: Anyone could do it. But we've got to make them do it.

TOM: They don't want to spend the money rewriting all that stuff. My credit union may not have a fancy iPhone app, but they allow non-alphanumeric characters in my password. So that's why I like a credit union.

Steve: Yeah. And he says Chase banking doesn't. But a credit union, some small credit union often does.

TOM: Because they don't have that legacy. That's the reason. Question #8, a listener requesting anonymity, or however you pronounce that, shares the "inconvenient truth" about Apple's app vetting. Here we go. He says: Every so often I hear you talking about how the apps in the Apple App Store are somehow more secure because they've been vetted by Apple. On first glance, this appears to be true. Apple does some vetting for each app. However, the truth is a little different. It's trivial to get undesirable code past the vetters.

As an active iOS app developer, I thought I'd share some insider information: For example, in my apps I allow four weeks to pass from the time of submission before doing anything that might be regarded as nefarious. I even check the date at time.nist.gov to be sure I'm picking up the real date. I host a website with a simple text file on it. This contains instructions for the app, allowing me adjust the app's activity. In my case, I use this to turn user logging off once I have enough data, but it doesn't take a genius to work out how this could be used to activate a malicious payload.

In my case I use this data to improve the apps and see which features are being used. I've got six apps in the App Store, every one of which sends data back to me behind Apple's back, without Apple or my apps' users knowing a thing about it. It's trivial to do, and there's almost zero chance of being caught. Great show, been a listener since the start.

Steve: So there's a perfect example of what we've talked about often, which is there is no way for Apple to know exactly what an app is doing. I mean, they'd have to have the source code, and then have to go through and inspect the source code in detail to see what's going on. So this anonymous listener, who is a developer, obviously came up with a slick way around it. He has his apps look for the date and change their behavior after a certain date. So it gets past Apple, who checks the app to see what it's doing at time of submission; and then simply, since the app knows it's going to have network connectivity, every so often it checks the date. And if it's been long enough, it suddenly awakens an aspect of the app that was lying dormant before. And in this case the app pings this guy's website to obtain updated instructions about how it should behave, whether or not it should still log and where it should send the logs and so forth.

So, I mean, I wanted to share this because this is actually happening. And it's very clear there's just no way to prevent this kind of behavior. And as he says, his use is not malicious, but it's certainly the case that it could be. And so the point that I made last week with Leo, when we were talking about what's more secure, Apple or...

TOM: Android.

Steve: ...Android, exactly, is none of the above. Maybe Apple is putting a little more oversight on - we made the point that Android developers pay $25, and it's easy to be anonymous, and then they were able to dump apps day and night into the Android store. Leo countered with the fact that, yeah, sure, but those apps can also be removed retroactively. So both Apple and, for example, Google in the case of Android Marketplace, have the ability of pulling things that are later found to be a problem.

And my point is, install as few things as possible. Or look carefully at the reputation of the companies whose apps you're installing. Now, this guy, who's a multi-iOS app developer, he's doing something that Apple doesn't officially approve of because he feels Apple's policies don't give him the flexibility he needs to deliver the best app to the users. I wouldn't disagree with that. And there is no way, I mean, this is a perfect example of there is no way for Apple to guarantee the performance of their apps, no matter what they do. So any malware could do this. And it's not like we're letting any secrets out of the bag. I mean, this is clearly obvious to any developer who wants their app to work, who wants this kind of flexibility and freedom. There isn't any way to prevent it.

TOM: Any computer that has network access is going to be able to do this. That's just the way they work.

Steve: Yes. And these mobile devices are now where the bad guys are having their fun and having their jollies.

TOM: Chad in Omaha, Nebraska shares an unintended consequence of VPN. He says: Dear Steve, I wanted to share with you and fellow Security Now! listeners an issue I have come across. I recently placed an order through Roku. I ordered at my local open WiFi cafe while my connection was protected by VPN. But Roku cancelled my order and refunded the money. When I called to ask why, they explained that because the IP address of the computer I ordered from did not match the area of my credit card billing address, they flagged the order as fraud.

I explained at great length that the order was legitimate and that I was willing to reorder from my home computer so they could see that the billing ZIP code matched my home IP address location. But after a week of relentless calls, Roku still refuses to let me place my order again. Perhaps my situation could be a eye-opener to another Security Now! listener. Love listening to all the TWiT shows. However, without a doubt, Security Now! is top of my list. Thank you, Chad.

Steve: I thought that was really interesting. First of all, I am heartened from a fraud prevention standpoint to see that we're beginning to match up IP addresses with physical addresses. I mean, there's always been sort of an IP location technology, never worked very well. But over time, especially with things like Google roaming around, pulling the locations and mapping the locations of all of these WiFi nodes, we're beginning to get much more IP location granularity. And of course smartphones with GPS that also have IPs, that's helping to create a map of where, physically where, given IPs are located. And so the idea that that's now being used as fraud prevention I think is very nice because it means that people in Russia are going to have a much greater difficulty using credit cards from Omaha, Nebraska.

But there is an unintended consequence that a VPN provides. Because if you use a VPN, and this is related to that tracerouting example from a couple questions ago, your physical IP will not be where you're located. It will be where you're terminated. It'll be the other end of your VPN tunnel. And if you were using some third-party service like HotSpotVPN, for example, you're going to be connected somewhere probably remote from where you are, some distance away. And so if anyone tries to geolocate your IP address, they're going to see the other end of the VPN, not you. So on one hand I say, hey, Chad, nice going that you were in an open WiFi environment, smart about using a VPN to protect yourself. But, whoops, there was a side effect of that, which is you came out on the Internet with an IP of your VPN provider, not located near where you went into the VPN tunnel. It's a very cool problem.

TOM: Yeah. And hopefully companies like Roku will get to understand that, hey, if I call you and say, look, it's really me, there is a way to verify that over the phone and take the order. Come on. Don't just blacklist the guy from ever buying stuff.

Steve: It really does seem like they got a little carried away with this.

TOM: Yeah. Marcin Ceglarek in Gdynia, Poland - and I'm sure I've mispronounced your name, so I apologize - has a question about lithium ion battery management. He says: For some time I've been a happy smartphone user, and your advice about battery management has been really helpful. I feel much more confident about proper battery management now. Previously, I was discharging the battery all the way down, and then charging it to full again to avoid the memory effect from NiCad batteries, which I now know lithium ion batteries don't have and, in fact, is BAD for lithium ion.

But there is one more issue: Is it safe to leave the phone plugged in overnight? From experience, I know it will get fully charged in about three hours. Since I'm plugging it in around 8:00 p.m., it is fully charged by the time I go to sleep. But if I then unplug it, overnight it loses about 5 to 10 percent of its battery life, so in the morning I only have about 90 percent remaining. On the other hand, if I leave it plugged in all night, I have full 100 percent in the morning. What's the best approach?

Steve: Okay. So assuming that the phone or laptop is properly managing its lithium ion batteries - and we have to make that assumption. I mean, if you've got an older device which is causing batteries to catch fire, or doesn't have - I mean, as some have...

TOM: Yeah, yeah.

Steve: ...or isn't properly managing batteries, then we're turning you into the battery manager. And that's just a bad idea. What lithium ion, the way lithium ion batteries behave is different from the way NiCads behave. The way NiCad batteries are charged is a NiCad battery voltage will increase to whatever level it's going to and then begin to decrease. And so in fact there were, like, rapid NiCad chargers that the RC modelers used for a long time which could recharge a NiCad incredibly fast by watching for that dip, watching for the point where the voltage began to drop. And the second it was detected to be dropping, the NiCad battery charger would stop and say your battery is now fully charged.

Lithium ion doesn't work that way. Lithium ion has to be stopped charging per cell, that is, every individual cell in a series connected chain - remember that "battery" itself, the word means a multiple. You have a battery of guns. That's a bunch of guns. Or a battery of cells is what we refer to as a battery. So that's a set of cells, individual cells connected in series. Each cell has to be monitored separately, which is why, if you look at the connectors on our laptop batteries, you'll see sort of like a comb of connections. And when you put your battery in your laptop, that comb is mating to a comb on the underside of the laptop, and those individual connections give your laptop's battery management access to, that is, visibility into each connection between the cells in that battery module which you plugged in. Lithium ion has to be charged to a cutoff voltage and then stopped.

So what is happening in Marcin's case, he's asking, if it takes three hours to charge, do I need to unplug it at the end of that time, in which case it will then switch to battery operation and discharge 10 percent by morning, or can I leave it plugged in? The answer is we have to assume the battery management technology is going to do its job. That is, it's going to work correctly. It's the battery manager; you're not. So leave it plugged in. What'll happen is it will charge the battery, the individual cells in the battery, to their cutoff voltages and then stop at the proper point. And then those three hours or eight hours, rather, while he's sleeping, it will be running off of the AC, not off the battery. So it will charge it and then not drain it. It'll be running off AC so that in the morning he disconnects it, and it's fully charged because the battery was charged but has then been floating overnight, not been discharged by not being plugged into the wall overnight.

And so that's the right strategy. Trust the management technology. If the device is within the last five years, everyone has figured out how to do this right, and they are doing it right. So leaving them plugged in over the long term is fine, as long as you're going to be using it. And remember, do not discharge lithium ion any further than you need to. Try to plug it in as much as possible.

TOM: Yeah, that's what I do. I plug it in all the time. Eric Duckman in the chatroom pointed out you can turn it off. If you're not worried about taking calls overnight, while you're sleeping, that's the safest way is just don't have it on at all. And then it's not draining anything. But if it's not charged all the way, you probably want to plug it in so it charges up overnight. In which case you're fine.

We actually answered Brendan's question earlier on in the show. He wanted to know if you recommend the keyboardless Kindle or the Generation 3 keyboard version. Sounds like you like the keyboardless one.

Steve: Yes. The reason I put the question here, though, was that he was saying I want to buy one and don't know what to get. My feeling is it is impossible to recommend. It'd be like someone saying what's the best movie? Well...

TOM: "Casablanca."

Steve: What do you like?

TOM: Yeah, yeah, exactly.

Steve: It is too personal. I have all of the Kindles. I've shown them to various friends. Everyone likes a different one. Some people love the idea of touching the screen to change the page, even though the reviewers think it sort of sucks to do that. And other people want a physical button. They just like the idea of resting their hand on the button. I feel like, hey, I mean, since I've got them all, if I'm going to go down and walk on the beach, I'll put one in my pocket, I want it to be the small one. But if I want to, like, sit at Starbucks, I want something that's easier to hold, so I like having more margin at the bottom.

My point is, I really think they're all good. And so the one you choose is personal choice. It's what, you know, how sensitive are you to price because the smallest one is the cheapest one. If you get the ad-supported, it's only $79. The larger, older ones are more expensive. Maybe you care about touch, or maybe you don't. If you're going to be typing things in a lot, then having either the physical keyboard or the touch keyboard would make sense over having no keyboard and having to use the up-down left-right arrows to navigate around.

So at this point I think they're all good, and it's just a question of who are you? How sensitive are you to price? How sensitive are you to size? How sensitive are you to ease of handling? Do you want to have a physical button? I like physical buttons. Other people think it's cool to be able to touch the screen.

TOM: There you go. And that's good advice for helping figure out which one is right, once you've figured out who you are. But we can't figure out who you are.

Steve: Right. And there are now, I think, Kindle is carried by Staples, is it? Or Best Buy?

TOM: I know Target had Kindles.

Steve: Wait, no, I'm sure it's Best Buy because...

TOM: Best Buy has it, too. You're right.

Steve: Yes. Kindle Fires were being purchased at Best Buy. So really, I would say the good news now, even things that used to be, like Amazon, web only, are now in physical retail. I would say go to the store. Go to Best Buy. Hold them. Feel them. And that's the way to make your decision. Try not to do it online, if you don't have to.

TOM: Yeah. John Morton in New York City has our last question here. Actually it's a tip for Mac battery management. He says: I've been listening to you and Leo since day one of Security Now!, Steve. I had a quick recommendation for you regarding battery life since you've mentioned it several times in the past few episodes.

It's a Mac program - sorry, Windows users - called Watts from BinaryTricks.com. The program basically keeps track of your battery usage and prompts you through various stages of cycling your battery when the conditions it's watching for have occurred. It will replace your battery indicator in the top menu bar with a much more robust dropdown menu that reports the condition of your battery and other bits of info.

The feature that might interest you the most is on its "Notifications" screen. There are two very handy options under the "Long Term Storage" heading. The first is "Notify me when reaching 50% of battery charge" and the second is "Shut down MacBook when reaching 50% of charge." I know of no other way of getting your MacBook shut down at its optimal level of battery usage. I hope you find this little program useful. I have. And thanks for all the education you've given me since Episode 1.

Steve: So I took a look at it, and it looks very nice. It is not free. It's free, full function, for 30 days. And then the developer and author wants $6.95.

TOM: That's not terribly expensive. Is it in the Mac App Store? Or you just have to go get it yourself from them?

Steve: I don't know whether it is or not. I went to BinaryTricks.com, and this is the program that the developer is selling. And I looked at it. It looks very nice. And for someone who's, like, more interested in active participation with the status of his battery, it looks like it really does provide lots of information. It'll tell you the last time you recalibrated your battery and give you a popup when it's time to do a full-depth discharge and recharge in order to recalibrate it. As you read and he mentions, it'll help you store your battery, if you're not going to be using your battery and your MacBook for a long time. It's best to store lithium ions half charged rather than fully charged. It's best to use them in a full charge and recharge as quickly as possible. But if you're going to, like, stick it on a shelf for a few months, take it down halfway, and then that's better for long-term storage. So this was - I thought it was a great tip. And I have downloaded it and am using it on my laptop.

TOM: You can't get it at the App Store. I just checked. Although, remember, you're giving 30 percent of the profits to Apple in that case in exchange for the management of it through the App Store. So if you want Binary Tricks to get all the money, go to BinaryTricks.com.

Steve: Yeah.

TOM: All right. That's it. Thank you, Steve. This was great. Always a pleasure to be on Security Now!. Of course Leo will be back next week. But don't forget you can find all the goodness that Steve does at GRC.com, things like SpinRite, things like ShieldsUP. Anything new that you've got going on that you want to mention before we go?

Steve: Well, we've got Off The Grid, the Off The Grid paper-based encryption. The printing page is all finished. Still not linked to the main menu. I just have some more work I've got to get done on the other pages, just documentation. But that's where I'm working now. So hopefully I will be able to announce that it's available at the main menu. But you can just go GRC.com/offthegrid.htm, or actually I'll append the "htm" if you don't do it. So GRC.com/offthegrid, and that'll get you into the Off The Grid pages. And they're all interlinked from there.

TOM: I love that project. That's such a great project. Check it out: GRC.com/offthegrid. That's it for Security Now!. We'll see you next time.

Copyright (c) 2006 by Steve Gibson and Leo Laporte. SOME RIGHTS RESERVED

This work is licensed for the good of the Internet Community under the Creative Commons License v2.5. See the following Web page for details: http://creativecommons.org/licenses/by-nc-sa/2.5/