SECURITY NOW!

Transcript of Episode #334

Listener Feedback #134

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-334.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-334-lq.mp3

Leo Laporte: This is Security Now! with Steve Gibson, Episode 334, recorded January 4th, 2012: Your questions, Steve's answers, #134.

It's time for Security Now!, the show that covers your security and privacy online. And Steve Gibson is here. He is the man, the host, the legend, the guy who does this show every week with lots of great security news and information, explanations. He's a teacher, and he's a mentor, and he's a good old guy, which - one of my great friends.

Steve Gibson: Older and wiser, I hope.

Leo: Yes, let's hope.

Steve: First episode of 2012 for you and me.

Leo: Happy New Year.

Steve: Yeah, likewise, Leo.

Leo: We had a fun time last week. I have to say I'm glad we didn't do a rerun because we talked about sci-fi. That was so fun.

Steve: Well, it's something that's near and dear to both of our hearts, and it turns out it was a huge hit. I got a ton of feedback from our listeners who went to GRC.com/feedback to tell me what they were thinking. And also stuff in my Twitter stream, I mean, people just really, really liked it. In fact, I had one, I saw one note where the guy said, okay, my three favorite episodes of all time were Portable Dog Killer, the Vitamin D, and the Sci-Fi Special. And it's like, okay, wait a minute.

Leo: Not security.

Steve: None of those are about security.

Leo: I'm not into security, man. I don't dig security. I don't want to do security.

Steve: I think they just like hearing you and me ramble on about random things that...

Leo: You know, that's really what this network was originally is just a bunch of gas bags, talking about stuff. And we really probably should just stick with that format instead of trying to do content. Although we've got a lot of content. We've got a Q&A today. In fact, I bet you some of those questions have to do with last week's episode.

Steve: I avoided that because I figured, I mean, I had to literally step around all of the feedback about sci-fi because I thought, well, okay, we did that, but let's not - actually there were a lot of interesting comments.

There's a really interesting industry-wide problem that forced Microsoft to issue a very rare out-of-cycle patch, which they did, we'll talk about.

Leo: That's huge.

Steve: And some little bit of legal news. But, yeah, we have a - I did 12 Q&A questions that I found because a couple of them are just going to be quick comments and so forth. So, yeah, I think everyone's going to be happy with another podcast.

Leo: You know we only have really one commercial. So why don't we launch right into it. But before we do, just a program note. Next week is the Consumer Electronics Show, and as usual we're going there with our bags in hand.

Steve: Your batteries charged.

Leo: Batteries charged, yeah. I just got a ginormous battery for my phone. By the way, it comes with, as most batteries do, instructions for conditioning - it's a lithium-ion battery - completely in opposition to everything you have recently told us. It said

"Do not plug in your cell phone right away. Use it normally till the battery is completely discharged. Once discharged, then plug it into a wall charger, not a USB or a car charger. We want high-voltage, and then we want to charge it all the way. But don't do too much to drain the battery quickly. You want it to drain" - this is nuts - "for the first few times at a slow, even rate. Drain the battery fully again before plugging the phone in. Do not," it says, "do not fall into temptation to plug the cell phone in..."

Steve: Wait, wait, wait. Is this because of a - is this a larger capacity battery?

Leo: It is. It's about three times normal.

Steve: The reason they're telling you to do that is that it's necessary for you to train the phone...

Leo: Ah. It's not about the battery.

Steve: It's that you have to train the phone that, against its expectations, this battery is going to last much longer than it expects. And so...

Leo: I don't think that's what this is about, but I think you're right. That's probably a good idea.

Steve: No, no.

Leo: But they're saying things like "conditioned cell phone batteries are meant to be able to catch a few minutes of charge as needed, but it is the first few days that are vital to good conditioning." Which we know is not true; right? On lithium-ion.

Steve: Yeah, that's not true. Although it would be the case that your phone's battery...

Leo: The phone'll be confused, you're right.

Steve: Your battery meter wouldn't be expecting to have a battery lasting so much longer.

Leo: Right. That makes sense.

Steve: So you'd want to show it a few times, by going to the extremes, how long this particular battery lasts. And the phone's going to be going, whoa, okay. Because otherwise it would be showing you a very pessimistic-looking percentage meter, and you'd be saying, wait a minute, I got an extended capacity battery, and this thing's running at zero.

Leo: Good point.

Steve: Even though it wasn't even there yet.

Leo: That's, you know, I doubt that's what they're thinking. But I think that that's a good reason to do it. But just to recap what you said some time ago, a few episodes ago, there's no - lithium-ion batteries, the best thing to do is just keep them charged as much as possible all the time; right?

Steve: Yes. Now, the universal agreement is that, unlike prior technologies, prior chemistries, both nickel-cadmium and then nickel-metal hydride, which was actually the same fundamental electrochemistry, those had a memory effect. So if you only discharged them a little bit and recharged them often, they would forget that there was, like, a lot more room down below the point that you were normally starting the recharge. So the logic there was run them all the way down to the ground before you recharge them. And if you can't, then you could reset their memory by deliberately doing some deep cycles.

Lithium-ion, completely different story. It has no memory. And it actually is better for the chemistry not to be running it all the way down. So there the logic is charge often. And if you are in a situation where, for example, you have access to a plug, and you've got your adapter, always use the adapter when you can. And actually, one of the questions in today's Q&A explains a mystery that we encountered a few weeks ago. And I don't remember which episode you were on and which episode Tom was on with me. But it was someone's cell phone instructions said, once it's charged, unplug it from the wall.

Leo: All of them say that.

Steve: Which - okay.

Leo: And Nokia started it. And the reason is those little chargers draw power even when they're not doing anything.

Steve: That is exactly right.

Leo: So it's a green thing.

Steve: Exactly right. I was assuming that it's like, well, okay, the phone ought to be taking responsibility for battery management. And in fact the manufacturer is trying to take responsibility for those..

Leo: For the charger.

Steve: I like calling them "wall warts."

Leo: Wall warts, yeah. But we all leave our wall warts - c'mon, who's going to plug in and unplug all their wall warts all the time?

Steve: And it is annoying to feel how warm they are, which of course is the giveaway that this thing is expending electricity.

Leo: Make better wall warts.

Steve: And they could, Leo. They could easily have smart wall warts that shut down when - after they've done their charging cycle.

Leo: It just baffles me. And this is a company that makes these batteries. And they just - it says "Continue to drain your cell phone battery completely before plugging it in at every available opportunity. Regardless of the age or conditioning of your cell phone battery, it will always last longer over time when discharged fully before charging it again." Now, of course, there's circuitry to keep it from discharging fully. But it's wrong.

Steve: And you know, Leo, this, to me, this sounds like language that has not changed since they were...

Leo: Right. It's NiCad.

Steve: Exactly. So they've just - they've never updated. It's probably an Asian manufacturer.

Leo: It is, of course, yeah.

Steve: That paid for a good English translation.

Leo: Oh, that's what it is. This is like Colonial British. You know, when you're using nickel-cadmium batteries, 'tis always better to discharge fully. That's probably what it is.

Steve: Yeah. And they...

Leo: They just didn't want to buy another translation.

Steve: Yeah. The electrochemistry completely changed out from under them, and they never updated their manual.

Leo: That is so funny. Well, I'm just glad I listen to Security Now! so I didn't have to do that.

Steve: Yeah.

Leo: But I think you are right. I think it's probably a good idea to let the cell phone know. I wonder how smart Android software is about all that, whether it needs to - whether it'll figure that out, or how long it'll take to figure it out and all that.

Steve: It really, see, and this is it, it really can't...

Leo: It doesn't know.

Steve: ...because lithium-ion cells hold their voltage with a very slow drop until very near the end, when it suddenly drops off. So the reason our laptops sometimes do put us through a calibration cycle is they need to see where the end of the battery's life is in order to reset this timer. So they actually go on time more than on voltage because voltage isn't a reliable source of feedback for them about the current state of the battery's charge. And that means that you really would have to teach your phone that, against all logic, this battery is lasting three times longer than, like, the normal battery.

Leo: That's very interesting. Well, I'm going to do that. I'm going to let it die a couple of times and see what happens.

Steve: Cool.

Leo: So let us, well, you know what, let's do - oh, I know what I was saying. I was going to talk about CES. I got completely distracted. One of the reasons I got this big battery is so I can use this thing all day at CES. Most cell phones, I plug them in all the time, but CES is kind of an unusual situation where you're on your feet 12, 14, 16 hours a day, and you want these phones to last. A number of manufacturers have given me external batteries, backpacks with built-in batteries. They know. This is a universal problem.

So we're going to move Security Now! next week because we'll be on the show floor. Little difficult for us to do this kind of show there. So Monday, 9:00 a.m. Pacific, 12 noon Eastern, this coming Monday, January 9th, that's the next Security Now!. I'll remind you at the end of the show. And we'll do a show then. And then do watch, though, all week long for the latest from the Consumer Electronics Show. It's going to be a lot of fun, I think a lot of interesting things, yeah.

Steve: Yeah, cool.

Leo: All right. Let's get to our security news, Mr. G.

Steve: Okay, now, this is really interesting. Back in 2003, some researchers raised everyone's awareness of some fundamental problems with hashing algorithms which were being used in all of the server-side tools. And by "all," I mean just across the board, in Oracle's Java and Glassfish, Microsoft's ASP.NET system, Python, Ruby, PHP both versions 4 and 5, Apache's Tomcat and Geronimo, Jetty, Plone, CRuby, JRuby, Rubinius, and even more recently the v8 JavaScript engine.

What happens is, when users are submitting data from their web browsers to a web server, and we've talked about this in all kinds of different contexts, a so-called "post query" is made to the server. And the body of the query contains all the arguments that you're sending. So when we're filling out a form during purchasing, on a web forum where we've got multiple fields, any time we're sending data up to the server, we're using almost always a post-style query containing data.

Well, what the servers do is they hash the various elements of the request as a means of storing it in a data structure that then allow these server-side scripting languages, thus Python, Ruby, PHP, Java, ASP.NET and so forth, they create this data structure which those scripting languages then query in order to obtain the arguments that the user has submitted. Well, back in '03, so, what...

Leo: Eight years, nine years, yeah.

Steve: Yeah, like eight or nine years ago, plenty of time ago, researchers pointed out the fact that it would be possible to create, to deliberately create hash collisions which would overtax servers. Now, the reason this happens is, if you - we have to explain a little bit about how these hashes are used. The idea is you sort of use the hash, the result of the hash, like an index into a data structure where you store stuff. So, for example, say that the user submits some data and individual pieces are hashed down to a 32-bit token. We're not talking like crypto hashes, which are 128 bits or 160 bits or 256. We're just sort of wanting to distribute the incoming data evenly in a bunch of separate little buckets. So we might turn the input data in individual pieces into sort of a quick and easy, like a 16- or a 32-bit hash. So that's then used sort of as the index to store this piece of data, the idea being that all the data that you hash is going to sort of have different 32-bit values, and so it's stored under different 32-bit tokens.

Well, these researchers in '03 said, okay, but if they weren't stored in different 32-bit tokens, if they were all stored in the same one, like by somebody maliciously designing the data so that it would hash to the same value, and since these are not cryptographic-strength hashes, they're just sort of little quickie hashes meant to distribute the data evenly, that means that it's possible to pretty easily design them so that they don't, that is, design the input data so that it isn't distributed evenly, you can force all of the input to hash to the same 32-bit value.

Well, it turns out that our servers can't handle that at all. I mean, the idea is that this approach for average typical data is very quick and does a very good job. But the worst case is very bad. It collapses. So what happened was that several months before reminding everyone of this in the 28th Chaos Communication Congress, which occurred recently, I think it was, like, last Wednesday or the Wednesday before in Berlin, security researchers who decided they wanted to demonstrate this, they notified a couple months before all of the various makers of the server-side technologies that I talked about, I mean, like all of them that were vulnerable, Microsoft included, and all the other guys. And they said, hey, we want to let you know, this has been known about for eight or nine years. You really need to fix it because we're going to...

Leo: That's a long time.

Steve: Yes. We're going to show people how bad this is. Okay. So how bad is this? They demonstrated that a single person with just using 70 to 100 kbits of connection bandwidth, which we know is like nothing, could completely saturate a server using an Intel i7 core processor.

Leo: A single machine.

Steve: A single machine could be brought down, essentially creating a denial of service attack on a website with 70 to 100 kbits of queries. And what they showed was that hours of CPU time can be consumed making a single HTTP post request. So they, in one of the papers which they demonstrate, they show that, if a single attacker had a gig connection, that gigabyte connection, okay, or if an attacker had 10 100MB connections, like from a small bot fleet, or a hundred 10MB connections, which is easy to get, that could fully saturate 10,000 i7 core servers.

I mean, a huge commercial website can be taken down with relatively low bandwidth usage by essentially perpetrating a CPU resource consumption attack by generating these requests to the server which deliberately cause hash collisions. The idea is that these hash collisions prevent the data from being scattered out over space, essentially, where the data can be accepted quickly, and force all the data to be stored in the same place, which is like a worse-case scenario over at the server end for it accepting the data.

So Microsoft a couple weeks ago issued an out-of-cycle patch, and everybody else has in the meantime fixed it. So the problem is, now that this is known, the reason Microsoft was of course motivated to fix this, is it's trivial to do this. It's no effort at all. And the industry could easily see that bad guys would just have fun doing this. I mean...

Leo: But just so people understand, this isn't a threat to your system being compromised or anything like that. It just means that bad guys can do DDoS attacks, attacks on websites easier, with less hardware.

Steve: Well, actually it's with almost no bandwidth.

Leo: Very easy. A gigabit. A gigabit's not "no" bandwidth. That's a lot of bandwidth.

Steve: No, no, no. No. But a smaller server, I mean a state-of-the-art i7 core server, can be DDoSed, actually just DoSed, you don't need distributed, just DoSed, with between 70 and 100K.

Leo: Oh, I get it.

Steve: Yeah.

Leo: So if you had - I see. Wow.

Steve: Yeah. So, I mean, so the idea would be you'd need much more to bring down a 10,000-server farm, but even that can be brought down. But a single server, you can take it down with 70K of data, tying up its processor for hours, making a single query. So this is - we've talked about how the generic concept of a denial of service is that regular users cannot access the resources of a web server for some reason. Well, there are many different reasons. One was like making lots of connections that you don't ever complete. And so you exhaust the server's ability to accept valid connections by tying it up, waiting for bogus connections to actually be made. So that was a TCP connection resource denial of service where you just - you used up the server's ability to accept connections by sort of hanging connections that weren't actually made.

A different approach is just flooding it. A bandwidth denial of service is where you just pound the actual connection to the server. The server's sitting around with nothing to do because you've, like, brought the routers down. The routers can't handle all the traffic trying to get to the server. So valid users are unable to get there. So that's typically a distributed denial of service attack. The outcome is the same: valid users can't get there.

So what we have now, this thing is a new type of denial of service where, by cleverly designing the queries that are made to the server, we are overloading the computational ability. And so we have a low-bandwidth, just a trickle of queries can take a server off the 'Net because the way it's accepting the data will collapse if the data is designed just right. And it turns out it's no rocket science to do this.

Now, all you need to do, the reason hackers can do this, is all of these different systems that I've talked about - Java, ASP.NET, Python, Ruby, PHP, Tomcat and so forth - they use a very simple and well-known hash. There's only a few of them that most people use. It turns out it was Dan Bernstein's hash, variations of a hash that Dan Bernstein came up with to be very efficient in computation time so that it wouldn't take much time to do this. But that's what's been exploited.

So Perl is one of the systems that took this seriously in '03. They're not among the list of vulnerables, they haven't been, because all they had to do was just you randomize the hash when you boot Perl up, and so each instance of Perl running all over the Internet has a slightly different hash, so there's no way for the bad guys to predesign data that will cause a deliberate collision of the output of the hash. That's all it took. But nobody else bothered until these guys said, okay, it's been eight or nine years, let's make everybody bother because, without fixing this, all of the servers out there were vulnerable until just recently.

So this was a good thing. And that's why Microsoft, if our listeners noted, issued an out-of-cycle patch. They fixed three other things in ASP.NET, as well. One actually was a critical security vulnerability that they fixed along the way. But what really motivated them to do it was that they realized any IIS server could be taken down just by making some tricky queries. Really, really interesting, yeah.

Leo: Wow, very interesting.

Steve: Now, this one is obscure, but I know that the VLC video player, which is one I often use, the VideoLAN player, is extremely popular. So I wanted to make note that a vulnerability has been found.

Leo: Oh, yeah, because we all use that. We recommend it heavily, yeah.

Steve: Yes. A vulnerability was found in an obscure file format. It's actually the TiVo file parser has a vulnerability such that, if you were to open a .ty file, that's the file format that TiVos use, the TY file demultiplexer plug-in, which is libty_plugin.*, there's a known vulnerability in that that allows a remote code execution. So what would happen is somebody would send you email that contained a small file, a .ty file embedded, saying click here if you want to see a funny movie or something. Or you'd go to a website that would, in one way or the other, induce your system to play that. So you'd have to have VLC installed. Versions 0.9.0 through 1.1.12 are vulnerable. So only the latest 1.1.13 fixes this problem.

So I just wanted to give everybody a heads-up. I, in looking at the VLC site, it didn't look to me like this 1.1.13 was available yet. So what you can do, and in fact I don't know how many people are watching their TiVo files in the raw like this using VideoLAN, but you can simply delete the libty_plugin.* files from your VideoLAN plug-ins directory, and then you're safe, too. So if you're not using these .ty files, I would just say get rid of the libty_plugin files.

Leo: And what do you use them for again? What is the...

Steve: Actually it's only for playing TiVo files.

Leo: Oh, okay. Delete it.

Steve: Of course everybody - yeah, exactly, delete it. Everybody gets it as part of the package, so that if you happen to run across a TiVo file - and the problem is it will try. If someone sends you email with a .ty file in it, and you're a VLC user, they could take over your computer remotely. So I just thought that was worth telling our listeners about.

Leo: Yeah, [indiscernible], we all use it.

Steve: Yeah. I do. I've got it on all my machines. It's a great player.

So in privacy news, the Ninth U.S. Circuit Court of Appeals has upheld the constitutionality of 2008's FISA, the Foreign Intelligence Surveillance Act, which had come under fire. What happened was, and we've talked about this in the past when this was just new and happening, there were 33 different lawsuits which users of telecommunications companies, including for example AT&T, Sprint Nextel, Verizon, BellSouth and so forth, when it became known that the NSA was asking those companies to surveil U.S. citizens, a bunch of lawsuits were filed against those telecommunications providers, saying we're not happy that you're surveilling us, so we don't think that's right. And what just happened, after many years, is that it percolated up through the legal system. A lower court said yes, in fact, this is constitutional, so these 33 suits should be dropped. And so that decision was appealed to the Ninth U.S. Circuit Court, that ended up upholding that lower court ruling.

So the EFF, our friends at the Electronic Frontier Foundation, and the ACLU, the American Civil Liberties Union, are both unhappy that the lower court ruling was upheld. A judge, Judge Margaret McKeown, who was on the Circuit Court of Appeals, said that "Electronic intelligence gathering depends in great part on cooperation from private companies ... and that if litigation were allowed to proceed against persons allegedly assisting in such activities, the private sector might be unwilling to cooperate with lawful government requests in the future." So I would say she was right about that.

Leo: We kind of were hoping that would be the case.

Steve: Exactly. So unfortunately, it looks like the FISA has been ruled constitutional, and so those suits dropped.

Leo: And President Obama signed into law the National Defense Act just this past week which allows them to hold without...

Steve: Detain...

Leo: ...indefinitely without warrant. It actually overturns habeas corpus. But I guess that's not in the Constitution.

Steve: Yeah.

Leo: It's a little scary. We live in a - we're starting to live in a police state. I hate to say it.

Steve: So over the week between Christmas and New Year's, records were broken, Leo.

Leo: Yes?

Steve: Not old-fashioned vinyl records.

Leo: No, but it's about time.

Steve: Maybe those. 20 million Apple...

Leo: I saw that.

Steve: ...or Android phones were activated. And here's the big number that caught my attention: 1.2 billion apps were downloaded.

Leo: In one week.

Steve: One week.

Leo: It's the first billion-app week. But if you think - that makes sense because all those 20 million people have new phones, and they need to put apps on them.

Steve: Yeah, they've got to fill them up.

Leo: They've got to load the apps on.

Steve: Against all of my advice. And it's like, yeah. So how many apps would that be per phone?

Leo: I don't even want to do the math. It's a lot.

Steve: Yeah, about a thousand.

Leo: I would say that most people who have smartphones have fewer than a hundred apps. We did a survey once. People like me have more than a hundred apps. But a hundred seems to be a cutoff point for, like, super geeks. Most people have. In fact, I look at most people, and they have a half dozen or a dozen apps. They don't...

Steve: Well, yeah, and thank goodness for folders, at least under iOS, because, wow, I mean, otherwise you're just, like, scrolling forever. It's like, where is that thing?

Leo: Yeah, they really needed that on iOS, for sure.

Steve: That's a good thing.

Leo: Yes.

Steve: And so also I ran across this very funny little note when I was going through the mailbag today for today's Q&A. And I got a kick out of it because it reminded me of something I was sure I'd seen once before, where RPN, that is to say Reverse Polish Notation, was regarded as "Yoda Speak."

Leo: It is. That's how Yoda speaks, in RPN.

Steve: Exactly. And I loved - I saw, "Impossible to see, the future is." Anyway...

Leo: The verb is at the end.

Steve: Exactly. Noel in Melbourne, Australia, he wrote - and I paraphrased a little bit because his note was longer and rambling. He said, "I was so stoked to learn from you of HP's revival of their venerable HP-15C calculator," which I notified people of, and again, I don't remember if that was with you or with Tom, Leo. But...

Leo: No, it was with Tom. I didn't know that. That's fantastic.

Steve: It really is. That is a - I still have a bunch of 11Cs. Now I have a couple 15Cs that are brand new as a consequence of learning of this. Anyway, continuing, Noel said, "Calculators should be landscape, not portrait. And of course RPN. But my wife is not a fan and won't touch my HP calculator at all. She calls my HP RPN calculators 'Yoda,' saying '6 enter 9, multiply you will.'"

Leo: Well, you know, that's a good way to learn RPN because then it can all make sense. It's Yoda speak.

Steve: That's right. "6 enter 9, multiply you will."

Leo: Give it the operands, then the operator.

Steve: "And she wants to know, what's wrong with '6 times 9 equals' that we learned in school?"

Leo: We could go on and on with what's wrong with that.

Steve: And it's funny because I put "RPN" and "Yoda" into Google, and sure enough. I knew I'd seen the notion of RPN as Yoda somewhere before, and I got a bunch of hits on that.

Leo: So funny.

Steve: Anyway, but thank you, Noel for reminding me. I got a kick out of that. And speaking of getting a kick out of things, I got a nice note, actually today or, no, Tuesday the 3rd, yesterday, from a David Goldenberg, who said, "Hi, Steve and friends." He said, "I've been a happy owner of SpinRite for a few years now, and it's my secret weapon in the technology trenches. I'm the family tech guy. I help out at my kids' school and have my own part-time business fixing PCs, training, and networking. SpinRite is always within reach and never lets me down. Last week" - oh, you'll like this, Leo, I forgot. He also signed off with his call letters.

He said, "Last week I had been preparing a laptop for a presentation for my ARES amateur radio group. I volunteered to get a new program, NBEMS, which is the Narrow Band Emergency Messaging Software, running that sends text messages and email-type communications over the radio. After several days I had everything working great and spent hours getting screenshots for the PowerPoint I was to prepare. One morning I went to start up the laptop. As I was getting together with another ham to go over what I had and to get his machine working, I started up my laptop and got the dreaded BSOD" - of course the Blue Screen of Death - "and an unmountable boot volume error. I did not break a sweat or even worry as I knew from experience that SpinRite would save the day. And needless to say, two hours later, the drive was scanned, several sectors were repaired, and the laptop booted, and everything I needed was ready to go. You're great, and so is SpinRite. Thanks. David Goldenberg, KJ6MCQ."

Leo: KJ6MCQ, nice call sign.

Steve: Yeah.

Leo: Not W6TWT, but nice call sign.

Steve: Not quite, yeah.

Leo: Hey, we're going to take a break. We have questions for you from your devoted listeners who of course, as always, come up with lots of stuff, comments, a dozen questions coming in a moment. Are you ready, Steve? I've got questions.

Steve: You betcha.

Leo: Lots of them, starting with Dean Severson in Clearwater, Minnesota. He's wondering if there's a quick read about SOPA for non-techies. We've talked a lot about the Stop Online Piracy Act. He says: SOPA, MPAA, RIAA. As a geek around the Christmas table, I get the same questions as other geeks. Lots of the answers, which should be straightforward tech answers, really revolve around the battle with lobbyists. I tried explaining all this SOPA stuff to my friends and family. I get lots of blank stares: "What is he ranting about this time?" Is there a quick-read link you could send us to help the layperson understand what we're getting at when we try to explain these totally inexplicable issues? Thanks. Listener since Episode 1.

Steve: Yeah. In general, I don't know of anything. But for SOPA there is a bunch of good stuff at the EFF, not surprisingly. They're on top of these sorts of censorship and privacy concerns and so forth. And when I saw Dean's question, I just went www.eff.org, and their search system allows you, over on the left-hand side, to specify categories. And so I just chose the blog that they have and searched for the acronym "SOPA" within their blog. And it immediately found some very nice and written - I guess I wouldn't say to Dean that this is maybe for his family and friends, but it's definitely for him - that would give him a sort of a nice-looking sense for a way to describe these issues. So certainly the EFF is where I would go for things like SOPA that are lobby-based and privacy and...

Leo: Oh, yeah, they're the greatest.

Steve: ...and censorship concerns.

Leo: There is a website dedicated specifically to SOPA and PIPA, the two bills that are in front of Congress right now. It's AmericanCensorship.org, and they do have an infographic on this site that's very easy to understand, and videos you can embed, that you can show people or send people. There's a really simple video that you can email or download. There's a very nice infographic that talks about the number of people who participated in American Censorship Day. 6,000 people signed up. A million emails to Congress. You know, we kind of won that battle. I think a lot of the people who supported SOPA, including GoDaddy and others, realized this was something that the Internet did not want. But the battle is far from over. They will be back again and again and again...

Steve: I did also...

Leo: ...to try to break the Internet because the Internet fundamentally, to the Motion Picture Association and the recording industry, is a threat. They see the Internet and computers as piracy tools, not as anything else. And in order to preserve their business model, they want to literally break the Internet because they see it as dangerous. And we don't want them to break the Internet. We want them to change their business model to suit modern times. We can't go back to the 1950s, sorry.

Steve: Somewhere, just this morning, I forgot to follow up on it, because I was curious, but I saw a blurb that said that Hollywood itself is now putting together or has put together a petition or some effort of some sort to tell the MPAA to lay off.

Leo: Right. Well, yeah.

Steve: So even Hollywood...

Leo: None of this is monolithic. I mean, there are lots of young people in the record industry, in the movie industry, in the television industry who understand that, you know - it's not really them. It's the people who have the existing business model that don't want to change it that are doing this. And unfortunately the members of Congress don't understand technology, and so they don't know what's being asked of them and what the cost of it would be, et cetera, et cetera. Fortunately, I think that we kind of did scare them a little bit about this, and they're starting to think a little bit.

Steve: Yeah.

Leo: But we've got to keep on. We can't - the problem is that these guys will not give up until they win or they're defeated entirely. And I don't know how we defeat them entirely. We have the best Congress money can buy.

Steve: Well, and unfortunately when Chris Dodd went over there, I'm thinking, oh, no, because he's well connected and of course going to be a big lobbyist.

Leo: Yeah. Watch the video at AmericanCensorship.org. It's a very good place to learn about this.

Steve: Cool.

Leo: And share with your friends. Question #2. I'm glad you brought that up, though, because it's not over. It's far from it. We've got a long way to go. Mark Cykowski suggests that you stand on your head.

Steve: What?

Leo: What? Well, you wrote that. Steve, long-time listener to your Security Now! podcast. A few weeks ago you mentioned you returned your Kindle Fires in part because the on/off button got in the way, and you'd accidentally turn it off while holding it. It's right at the bottom there. And if you hold it - I don't hold it that way. I hold it on the edges. But it's true, it's too easy to switch. I agree. I was having the same difficulties until I found a solution. I just turn the Kindle around. Because the Kindle rotates, and the on/off button then is at the top. No, I didn't think of that. No more accidentally turning the screen off.

The only problem I've seen is that the opening screen and the shutdown screen messages are upside down. But I could live with that. I'm not having any of the problems you mentioned. And you may want to take another look at the Kindle Fire. I don't know if that's a solution. Thanks to you and Leo for all the great podcasts. Sincerely, Mark. I think that's prob- that's the "You're holding it wrong" answer; right?

Steve: Well, okay. So a couple things. Yes, of course he's right. You can...

Leo: You can do that.

Steve: The Kindle screen rotates easily. It also occurred to me since then that, if you had any sort of a case around it, and I would imagine everyone who gets a Kindle, I mean, it's fragile. In the same way that you're going to put some sort of a bumper case on your smartphone, I would imagine you would do that with a Kindle. And unless they exacerbate the problem by putting a big rubber bump there to sort of, like, make pushing the button easy, that might end up recessing the button so that, if they just had like a hole through the case, then you'd be pushing in in order to get to it. So that could solve the problem. I did want to mention, though, that I did get another Kindle.

Leo: Oh. Another Fire, or another regular Kindle?

Steve: No, a Kindle Fire. I got a Fire because I thought, well, first of all, it's so inexpensive, it's hard not to have one; and that I ought to have it in order to play with it and get to know it and see what I think. And now my complaint is that the aspect ratio is wrong for reading.

Leo: Oh, that's interesting.

Steve: It's right for widescreen movies, which is the reason it is widescreen. But I just don't like reading in a long column or in a really wide screen with very few lines. So it's like, eh, I just - I don't - I think they - I'm not a fan of the Fire. But I do have one.

Leo: I just bought another of the $79 Kindles. To me, these are the - this is the Kindle. And I got a case with a light in it. And it's not for me, it's for my wife because she lost it on the airplane, following in my footsteps.

Steve: Ah.

Leo: And so, but at 79 bucks, that's not nearly as painful as it used to be.

Steve: Yes. And that's just - that's a cute little one. Now...

Leo: I think it's the perfect Kindle.

Steve: I stick it in my pocket when I'm going to go somewhere where I don't want to be lugging something around. And what's really neat, too, is - you probably noticed this, Leo. On the back are two little silver contacts, which is the way the illuminated case gets powered by the Kindle. So at least the case that I've seen doesn't use its own battery. It borrows the Kindle's battery, which is very elegant because that means you only have one thing to charge, also.

Leo: Absolutely. Okay, let's - I've lost my browser. There it is. Joshua Gardner in San Antonio, Texas, says Firefox's memory usage has been fixed in version 9.01: A quick note, I switched to Chrome for a bit because I could not handle Firefox's memory management leakage, which we've talked about several times before on the show He says: 2.8GB of RAM usage with 30 tabs open. My poor laptop only has 4GB of RAM, so Firefox would constantly cause the hard drive to be running because it would have to swap things back and forth. However, I couldn't handle not having tree-style tabs, which you've recommended, the side tabs. So I returned to Firefox - because it's an extension for Firefox.

Over the last week I've had Firefox open probably 10 days it's been open with 30 tabs open all this time. And my current RAM usage is 200MB in the most recent version, 9.01. If you have a spare machine sitting around, you might want to run the current version with a few tabs open for a night or two, see if you get better results. I know you aren't a latest version kind of guy. Neither am I, in some cases. But in this case I think the benefit is there.

Speaking of which, what is with this new trend to get rid of the menu bar in software? That File, Edit, View, Tools, et cetera? I first saw it in MS Office. Yeah, they're using that ribbon thing now. That's their thing. And I'm seeing it in browsers. Now, yeah, they were talking about putting ribbons on Firefox. I don't - did they? And now a few other utilities I've seen have been updating in this direction. For the record, I don't like it. Thanks for the show and all you do. Josh in San Antonio, Texas.

Steve: Okay. So a couple things to note here. First of all, the phrase "My poor laptop only has 4GB of RAM." Okay, what is wrong with that picture? 4GB.

Leo: I know, it's a ton.

Steve: Oh, Leo.

Leo: Should be enough to last forever.

Steve: I mean, mine do, too, Josh. I'm not poking fun at you. I'm just saying our world has exploded. It's just, I mean, 4 billion. Billion. 4 billion bytes of RAM. Oh, my goodness, it's wrong. It's very wrong. So I also wanted to note, he mentioned v9.01 and wanted to make sure people knew that 9 was quickly replaced after its first day on the planet. But I noted something else just today. And that is, Mozilla is having problems - speaking of not having enough RAM - Mozilla is having problems because they are no longer able to build Firefox.

Leo: What?

Steve: They are having to back out of features because it will no longer build. They ran across a problem building in Windows, and so they added the /3gb switch, which reduces the - it used to be that Windows...

Leo: You're saying the memory footprint is so huge in Firefox that there's not enough memory on modern machines to compile it.

Steve: They can't compile it. Correct.

Leo: Well, I think they need to rewrite the thing, if that's - that's ridiculous. I can't even imagine.

Steve: I know.

Leo: If you can compile Windows, why can't you compile Firefox?

Steve: Okay. So back when Microsoft was architecting Windows - and remember that Bill Gates famously said that they had 10 times more memory than the Apple II.

Leo: Right.

Steve: The Apple II could go to 64K, and so they were going to go to 640K. Well, then they realized, okay, that's not such a good idea. We need more than that. So they went to a 32-bit platform. So they said, okay, 32 bits, come on, that's 4 billion bytes. We're never going to need anything close to that. So they arbitrarily divided the memory in half because why not? Neither half would ever get near full. So 2GB for the OS, 2GB for the applications. All the applications. Because, again, applications, how big could they be?

So at some point Mozilla hit their head on the 2GB size. And so they added what's called the /3gb switch, which is something you can add to the boot.ini file in Windows that moves that 2GB fence that divides the OS and the applications from the 2GB point up to 3GB. So the OS is squeezed down to 1, and that gives applications 3GB of space. Well, that's been hit now, too. So I was reading some of the developer blog comments in the forum apologizing for a couple people who were not on the email routing list because their code had been removed, and they weren't notified, because Firefox can no longer be compiled.

Leo: Wow.

Steve: So, yes, Leo, it is a...

Leo: That's stunning.

Steve: It is just a travesty. It's just, like, okay, guys, come on. This is just getting too big. And as you said, it needs to be started over. What they're going to end up having to do, apparently, and this is a big problem for them, they're going to have to switch to 64-bit OS builds to build the 32-bit version. And that's going to be, I mean, they're just - it's going to take them a while to do. But it just means we're talking about bloat with no end.

So, Josh, thank you for telling me that 9 seems to have solved this problem. I'm still happily on 3.6 right now. And at the rate those guys are going with Firefox, I mean, we do know that they're projected to have 12 ready at the end of April. Maybe this is going to throw a little kink in their calendar, since they can't compile it any longer. Ugh.

Leo: There's got to be more to this story. I - that's bizarre. Question #4, Jerry wonders what Steve uses and recommends, a hard drive that spins or a solid-state drive: After experiencing a hard drive failure many years ago I purchased SpinRite, now use it on a regular basis. I've never experienced a hard drive failure since. Here's my problem. I'm going to be purchasing a new computer, and I am torn on whether to get an SSD or an HDD. I want an SSD because of fast boot, fast startup of applications, silent operation, and SSDs generate less heat than HDDs. The two weaknesses of SSDs that give me major pause are no SpinRite support, and various issues concerning the successful implementation of full-disk encryption. I'd add a third, which is price, of course, price per capacity. So, Steve, what kind of disk are you purchasing for your new computers nowadays, and what are your comments, insights, recommendations on this dilemma I'm facing? Thanks.

Steve: I see a lot of questions like this in the mailbag which I sort of skip over just because I feel like, well, okay, this isn't the hard drive show. But I thought, okay. This is one that hits home.

Leo: So let's answer this once, and then you can refer to it forever and ever.

Steve: Yeah. I don't have any spinning media in any laptops.

Leo: Really?

Steve: Yeah. And I've got a bunch of laptops. My feeling is that's a place, just due to the delicacy of those 2.5-inch small laptop-style drives, that really scares me. So, that is, in terms of it's just - it's easy for the laptop to get bounced by mistake. You turn it off and shut it down. I mean, knowing what's going on in there, that this amazingly delicate little and ridiculously dense data stored on these little spinning platters, the heads are still flying, I, like, have to sort of sit there and wait until I'm sure that it's all stopped spinning and the heads have landed. And even then, you really do need to be careful with laptops. All of the experiences that the original hard drive iPod users had of their iPods dying is similar to what happens with laptops. And with laptops being so popular, it's just a problem.

So I do have large spinning hard drives in most of my big machines. But one of the first things I do is to take out the hard drive and exchange it with a solid-state drive for my laptops. I like my little Mac Air because it's just 64GB of solid state, and it just - it feels right to me that this thing I don't have to worry about. And the iPad is the same, for the same reason, that it's just - it's solid-state storage. It is just - that's not going to go wrong. But you're right, Leo, price is a concern.

And of course, for me, wear is a concern over on the SSD side. We've talked about how SSDs, the current technology does get fatigued when you write to it because essentially you're squirting electrons through an insulator to strand them out on a little piece of conductive island. And that electron tunneling, squirting the electrons through, essentially you're breaking down an insulator, forcing it to allow a charge to pass through. Well, that fatigues it. It actually does wear it a little bit.

So I'm in the process of building up a new server for GRC. It'll end up becoming GRC.com and www.grc.com and all that. It'll be GRC's new main server. And I've decided it's going to be SSD. But I did a couple things. I got highly over-provisioned SSDs. They have 28 percent over-provisioning so that the SSD has that much spare space which is available to be swapped in as it detects problems evolving. And it's in a RAID 6. So not only am I using the best, highest quality - and these are all - these are not MLC, by the way. These are all SLC. That's the other thing I do. And it is really expensive to go SLC.

But single-level storage is much more reliable than multi-level storage. Much more expensive because you're only able to store one bit in each cell rather than two or three or in some cases four bits. So it means that the technology is less dense, thus the SSD is much more expensive. But the reason is it's faster and it's more reliable. And so in addition to that, I'm not only using SSDs that are over-provisioned, but RAID 6, which means that essentially I have a four-drive RAID 6 array. So any two drives could fail, and the entire system still runs.

Leo: How much does that whole setup cost? You're talking five grand, easy.

Steve: That's about right, yeah. But I won't have to worry about it.

Leo: Ever.

Steve: So - ever, yeah.

Leo: Well, and I'll back you. I don't go that crazy. I'll back you, but I don't buy a computer now, even this big iMac, this 27-inch iMac right here that I use, it has a 256GB solid-state drive and a spinning drive for data. So it boots and runs apps off the solid-state drive. All my laptops now have solid-state drives. I don't do full-disk encryption. And I think that that's a good reason maybe to be suspect of those SSDs because as far as I know, I mean, we had Allyn Malventano - I asked Allyn, and he gave me a way to do it. But really it's not something you might want to do on an SSD, I guess, I gather.

Steve: Typically, laptops do support drive passwords.

Leo: Right, that's sufficient.

Steve: And all the SSDs support - yes. A drive password, I mean, most of the SSDs now will incorporate their own encryption. And so if you give the drive a password, then it's encrypting itself on the fly. So that's going to do the job.

Leo: That's sufficient. All right, good.

Steve: So, yeah, yeah. So...

Leo: Again, no, I love it. The speed difference is huge. It's well worth the price.

Steve: It is. Now, of course the problem is no one is going to afford a terabyte SSD.

Leo: No. Well, that's why I compromised. Now, they do make hybrid drives which Allyn, again, Allyn Malventano from PC Perspective is the SSD king, and I defer to him in every respect. He says those hybrid drives don't - they're kind of a mix - don't do as well. But I did the same thing, in effect, by having an SSD boot drive and a data drive, a terabyte data drive. Now I've got the best of both worlds.

Steve: Yes. And I hope you're swapping over on the hard drive, not on the SSD.

Leo: Yes.

Steve: Because that's what you want to do, too.

Leo: Right, right. Let's see. By the way, there's an interesting, according to aos101 in our chatroom, there's a Slashdot discussion about the fact that Firefox won't compile on a 32-bit linker anymore. So actually apparently it happened last year with 2GB. So they added a 3GB switch to the Windows build servers. And apparently it's not going to - it's good. If you want to read more about it, the developer section of Slashdot has a very complete discussion.

Steve: Incredible.

Leo: Yeah, it's really interesting. Really interesting. All right. Moving along to our next - now, every time I do that, I lose the question. Here we go. Jared. Is this the guy who wants me to talk in Australian? Western Australia? He wonders about...

Steve: No, he's Australian, but no. I think we have...

Leo: He doesn't want me to.

Steve: Believe me, we don't have an Australian who wants you to use your Australian accent.

Leo: Oh, oh, good. All right, that makes sense. It's not an Australian. All right, I get it now. Jared in Western Australia wonders about web browser referrer headers: Regarding the referrer header for a browser - by the way, he misspells "referrer," spelling it correctly with two - with three R's. It's so funny that it's misspelled in the...

Steve: In the spec.

Leo: ...in the spec. So as a result, all the libraries, everything have to use two R's, R-E-F-E-R-E-R, three R's, actually it's four R's. They don't double - they don't have a double R. Anyway, regarding the referrer header for a browser, Safari in this example, why are the headers grouped as they are? For example, some websites can work just fine on the iPad 2 with its larger screen area, 1024x768, same as many laptops. Gmail works great with full-blown navigation. But instead the website says, oh, you're on an iPad, I'm going to take you to the mobile domain. Based on this, it appears the referrer header is grouped so that iPhone, iPad, Mozilla are all grouped as one.

My thinking is, since iPhone/iPad are on the same line - see, he's misunderstood this completely - from the website's perspective, it's still a mobile browser, regardless as to whether it's capable of displaying a full-blown page or not, thereby eliminating user experience. While it's true some websites do give you a link to navigate to their full website, others don't. Is this a limitation of the referrer header somewhere in the chain? I don't mind the m. sites on an iPhone, as it is mobile suited, but an iPad? Is there any resolution, apart from using a desktop browser or hoping the web developer has linked you to their full-blown version?

Steve: So as you spotted, Leo, Jared's a little confused. It's not the referrer header, it's the user agent header. The user agent header has been around since the beginning. And we've discussed it in the past because it can also be a little bit of a privacy concern. Add-on things that you incorporate in your browser, like accessory packages or libraries that the browser has, can all tack their own version numbers onto the user agent so that every query your browser makes announces that this is the user agent that is sort of the client, the browser client, that is making the query. The logic, the concept was in the beginning that, if some user agents, for example, well, like really old ones would have been text only. They didn't handle graphics.

Leo: Links.

Steve: So, exactly, links. So the server could see what was making the query and then serve different content depending upon the requirements. It used to be that the user agent would also contain and state the resolution of the user's screen. Presumably, similarly, the server could then return content suited for the resolution.

Now, the reason Jared's question caught my eye is that I, too, and maybe you, too, Leo, have been annoyed when, for example, using an iPad, I'm given a much feature-stripped website which just doesn't do the same thing. And in fact, I have two different tokens registered for PayPal. I've got my original football, and I also have on my - I've got the VeriSign VIP on my BlackBerry, so that I'm able to use either. Well, what that means is that, when I'm using PayPal, and I log in, and they want my one-time password, they take me to a dropdown box where I choose which one. The problem is, the mobile version of PayPal doesn't offer that. So I'm unable to pay with my iPad because PayPal sees that I'm using an iPad, gives me their mobile version that doesn't offer me the option of specifying which one of the tokens I want to use. So it's like, argh, you know.

Leo: But that's not - it has nothing to do with how the browser's identifying itself. The browser's identifying itself. It's the website that's deciding what to do with the identification. So go to the - complain to the website, not to, I mean, it has nothing to do with the browser's agent string. The browser should always say what it is.

Steve: Oh, I completely agree. It's just it's an annoyance that the...

Leo: It's totally annoying, but that's the website's fault.

Steve: Correct.

Leo: It's stupid.

Steve: Although there have been instances where, I mean, there are add-ons, for example, that allow you, not for iOS...

Leo: You can change the user agent, yeah.

Steve: Exactly. And in doing that you masquerade what you're using, and you could tell the website, no, I'm not a mobile platform.

Leo: Right. Yeah. But that's - it seems like not a good solution, really. The website, you just send a note to the guy who did the mobile website. And I hope we don't do that. I'm pretty - I think I'm pretty careful when we do our websites to treat the iPad as a desktop browser. But maybe not.

Steve: And that would be the right thing, yeah.

Leo: Yeah. Because, I mean, it really is. It really is.

Steve: Yeah, it is.

Leo: Yeah. If you go to Leoville.com on an iPhone, you get this - I have a very nice mobile template. And the way it works, it's a WordPress site, the way WordPress works is it looks at the user agent, says oh, you're on an iPhone, good, I'm going to give you this nice mobile site. And I'm pretty sure, I'll have to check it on my iPad, but I'm pretty sure it does not do that on the iPad. It's smart enough to say, oh, no no no, you can see the whole thing.

Steve: Yeah.

Leo: Chris Wronski in Illinois asks us about alternative PDF readers: In Episode 332 you talked about alternative readers for PDF documents. I wonder if you and Leo would share your favorite choices? You first, Steve.

Steve: Well, now, I have to defer to you, Leo, because I purchased years ago a bunch of copies of, I'm not proud to say it, but Adobe's Acrobat, the full Acrobat system, the whole document preparation system and all that, which I've been moving forward and upgrading over time. So it incorporates along with it a plug-in for reading. And so I've got literally Acrobat Reader as opposed to just the regular PDF reader. So I haven't had to go looking for other stuff. So unfortunately I'm not a good resource for this. But I know that you are.

Leo: Oh, well, in some ways I'm not because on Macs of course you've got built-in software.

Steve: Built-in.

Leo: And there are third-party programs, like one of our sponsors, Smile Software, makes PDF Pen that lets you make them editable, fill out forms and stuff like that. So I have a series of tools that I use on the Mac, and I don't need anything from Adobe. And I'm very careful not to download anything from Adobe on the Mac. On the Windows side I use Foxit, which I like a lot. They make a distiller as well as reader. There's a free reader, but I pay for Foxit Phantom, which is the full thing. A lot of people like Cute PDF.

Steve: Yes.

Leo: Which is another really good one. There's one, I think it's free, called Nitro, that a lot of - so readers should be free. A viewer should be free because it's not creating PDFs. So Foxit Reader is free. Actually Chrome will read a PDF.

Steve: Oh, I knew they were going to. Has that been now...

Leo: Yeah, I believe so, yeah.

Steve: Okay.

Leo: So if you're using Chrome, you don't need anything. Let me see if there's any open source - I'm just going to look here and see if there's any open source for Windows. Because that would be nice. I bet there are. I bet there are. But I would say Foxit is my favorite. And I pay for the full Foxit because I think it's good. Although, ironically, they had similar security issues to Adobe's, which they fixed. But nobody's immune. No software is perfect; right?

Steve: Yup.

Leo: PDF, somebody - Dark in the chatroom is asking, is it an open format? I don't think it's an open format. But I think it's specified.

Steve: I think they did put - it's in the public domain.

Leo: It is, okay.

Steve: But I think it probably carries Adobe's copyright.

Leo: Something like - yeah, so it's a mix.

Steve: Yeah, yes.

Leo: If you don't mind using proprietary software, I like Foxit a lot. FoxitSoftware.com, and it's free. So I think that's a good choice. And they don't have - I'm pretty sure they don't have the switch that we hate so much that allows it to execute JavaScript and other programs within the browser. Or within the reader. Pretty sure that's missing.

Peter in Sydney, Australia solves the "unplug your phone once charged" mystery. Oh, we were talking about this.

Steve: And this actually is the one I was referring to, so we really have covered it. But we can just read it real quick.

Leo: Yeah. He says Nokia started this, and I think other phone companies now do it. I know my Samsung does it. The iPhone does not. Once it's charged, it says, in fact maybe all Android phones now do this, it says "It's charged. Unplug your charger." And he points out chargers consume power even if they're not plugged to a phone. And so that's why they're saying that. It's to save power. I don't - it's not much. But cumulatively, millions of wall warts all over the country, all over the world, that can add up.

Steve: And you know there are different technologies for chargers. The old-school black blob actually had a transformer in it, and you could put your hand on it...

Leo: It's hot.

Steve: ...and it's warm, yeah. I think that Apple's is a much newer technology switching charger. And so it may very well be that it is not drawing quiescent power when the device is not actually using it.

Leo: Switching, though, is much more expensive to implement. And that's, of course - figures Apple would do that.

Steve: And it makes them so tiny and cute and white and everything, yeah.

Leo: And white. All switching power supplies are white. Tom Burns in Chicago comments about Password Haystacks and Password Honeypots. He says, "If Leo is reading this on the air, I request his Australian accent as my favorite."

[Australian accent] Hey, Steve. Perhaps this is a bit obvious, but I thought it was worth mentioning. I'm not going to read the whole thing that way. We would lose our entire audience. First, no matter how long your password is, if it's the same across sites, then it's susceptible to password honeypots, sites that would capture your username and password either intentionally or through being compromised, and attempt to reuse the password and ID elsewhere.

So let's say you sign up for Twaller.com thinking it's Twitter or something, and you give it a login, a password, unique, but except that you use this same password everywhere. Then they would know and try your password elsewhere. Second, and perhaps this is a bit farfetched for the moment, any site that can capture your password and initiate a robo-login attempt with the same credentials across all of the common banking sites can be trained to look for low-entropy passwords and flag the padding for human investigation. Oh, that's interesting. If your technique of password padding became commonplace, this would be the next logical place for hackers to go. I don't think so because there's so much low-hanging fruit, they're not going to do any extra work.

Steve: Exactly.

Leo: Your listeners may want to either not use padding for less than completely trusted sites or have different types of padding for different types of sites. Tom Burns in Chicago.

Steve: So this notion of a password honeypot was interesting because I took it to mean something, a variation on things we've seen before. Remember once upon a time there were sites that offered you to sign up for contests. And, for example, we know...

Leo: Punch the Monkey and all that stuff, yeah.

Steve: And they would, yes, and they wanted your email address. And, I mean, you ended up getting spam as a consequence of that. So they were harvesting email addresses and telling you that, oh, sign up for this, and then we're entering you in a drawing, and there's a chance, I mean, we're talking old school. This was a long time ago. But a lot of people were doing that.

And so you could certainly imagine a site which wants you to create accounts, asks you to create an account where you identify yourself with your email address and a password. They're not saying use the same - they're not explicitly saying use the same password you use everywhere. But they're assuming that people are going to, if they're not very security conscious. And then they go and try to log on with the same credentials in lots of other common websites. And if someone is using the same email address and password, that's going to succeed.

So it certainly is possible for a malicious password honeypot site to be created. And I wanted to take this opportunity, one of the reasons I saw this question, is just to remind people that I've, with things like Password Haystacks and the one-time passwords and Off The Grid and all these things that I've done, those have been sort of research and experimental things. I'm using LastPass. I'm completely happy with LastPass.

Leo: Yes. And you're generating a fresh password each time.

Steve: Oh, absolutely. I've got something, I mean, I'm now dependent upon LastPass. Every so often I'll get worried about the fact that I don't know what any of my passwords are anymore.

Leo: Right. Me, too.

Steve: So I'll make a backup copy of the LastPass Vault offline so that I have it all. And sometimes something, if I don't have it, it's like, oh, shoot. And I'll have to go manually open up my LastPass Vault.

Leo: Right, turn on LastPass, yeah.

Steve: Yeah, exactly. I mean, but the good news is it saves us from these kinds of vulnerabilities.

Leo: Right.

Steve: So just wanted to make sure that people know that I'm still LastPass.

Leo: Yeah. And then I also use - the key is don't use the same password all the time. And I use a variety of ways to generate passwords for sites. Let's leave it at that.

Steve: Yeah.

Leo: Steve in Piscataway, New Jersey suggests that the "Router Reboot" problem is not so bad: Regarding the router rebooting which allows direct PC connection to the

Internet - I mustn't have been here when you discussed this.

Steve: Right.

Leo: Isn't it true that a PC that has already booted has a locally assigned IP address (e.g., 192.x.x.x)? If so, doesn't that mean that it is not routable from the Internet? Therefore there's no exposure, unless the PC happens to be rebooting at the same time the router is in "switch mode"? You're going to have to 'splain to me.

Steve: Yeah, it's interesting, Leo, it was something that came up, I think when you were in Paris.

Leo: Yeah, must have been.

Steve: And that is, it was revealed that some of the fancier routers that are doing a lot of things are, as we know, typically a Linux OS; but that the other layers of technology, like the NAT and the stateful packet inspection firewall and so forth, are additional services that do not start up immediately. And we heard it has been verified that some of these routers are simple bridges between the Internet and your internal network until they get fully booted, which means your home network is exposed to the Internet with no protection while the router is rebooting. Isn't that interesting?

Leo: Hmm. So...

Steve: And it is the case.

Leo: It is the case. Hmm.

Steve: Yes. So what Steve has said is, like, wait a minute, the machines on your home network would be private IPs, 192.168.what.what. So even if the router was rebooting, your machines would still have a private IP so they wouldn't be able to get on the Internet. What we realized a couple weeks ago was that, if your PC was asking for DHCP renewal, if you were renewing your DHCP lease, while the router was rebooting, you would actually - it would pass right through the router, and you would get a DHCP IP from your ISP, that is, a public IP rather than a private IP, because your router is just bridging the Internet traffic through.

So Steve's right that your machines on the private LAN would have private IPs. But the important thing is that their stack, their TCP/IP stack, would still be exposed directly. So it's not that they wouldn't be able to get on the Internet, which they wouldn't because they'd have a private IP, 192.168.whatever, but that incoming traffic would be passing directly through the router and hitting their machine. So the good news is, and we did talk about this at the time, most PCs have software firewalls now that are running, that are protecting them themselves. But on the other hand, we are depending upon our routers often for security. And so it is the case that the router is not providing us that unsolicited incoming packet dropping and security during this interval, until it finishes booting. So I would contend there is still some window of vulnerability which is just to sort of keep aware of. Yeah, interesting.

Leo: Anthony in Oregon wonders what to do when updates fail: Hi, Steve. I'm sure you know some of Microsoft's updates fail, for instance the latest one, KB2618444, the cumulative security update for IE 9. Even the support forum hasn't been much help. I do not use IE, but I want to keep it up to date, of course. Will my system not having this update adversely affect my security, even though I don't use Internet Explorer at all? And, if so, how can the updates that fail even after multiple tries be downloaded? I don't remember this subject specifically coming up on the Security Now! podcast. No, but it comes up all the time on The Tech Guy. If you think it's relevant, maybe you could address it briefly. I've been listening to Security Now! for years and enjoy it very

much. Keep up the good work.

Steve: And Leo, I don't know if you've got any magic solution.

Leo: Oh, I do, yeah. Comes up a lot.

Steve: Yeah. What do you tell people?

Leo: Well, there's two things. I'll answer both questions. First of all, updates, think about it, if you're doing an update in place of an operating system, it's like changing the table cloth by whipping it out and putting a new one in there. It's not always going to work. In fact, it's amazing it works as well as it does.

Steve: That was going to be my first comment, was I'm not surprised when it fails, I'm surprised when it doesn't.

Leo: I mean, really, you're modifying a system, you're building a plane while it's in flight. You're changing the engine, so - to add yet another metaphor. So when they fail, what happens often is - and this is actually the more serious issue. If an update has failed, it will not then go on to do other updates, and you will not be able to add future updates until that one update gets done. It's blocking.

So Microsoft has a very long tech note - I'll see if I can find it again, we put it in our show notes all the time at The Tech Guy show - explaining the 38 different things you can try if you've got a failed update, including clearing the registry. Mostly it's clearing a registry setting, undoing it and so forth. Now, if you have an update that you really want, you don't have to use automatic update, in-place update. You can change Windows Update. If you go into Windows Update and go into the settings on the left there, you can do a - because think about this. Administrators at a business with hundreds of PCs, they don't want to go to each PC one by one and do an auto update. They download the file once and then apply it over the network, or they go from machine to machine.

So Windows Update does allow you, you can actually go to the updates and download individual update files all by themselves. This is what I do, for instance, for system packs because they're so big. I just want to put the 700MB download on a USB key and then update all my various systems, yeah.

Steve: And carry it around, yup.

Leo: So you can go into the settings of - I don't remember the step-by-step, but it should be pretty apparent. Go into the settings and go to, I think it's called the "Catalog" of updates, and go to individual updates and get them. If it's stalled, though, I'll tell you before this show's over - we've got a couple more questions. I'll let you answer those, and I will find the knowledge base article. You can also Google - go to Microsoft. Don't even Google it. Microsoft has a great site, support.microsoft.com, where you can search for "failed update," and you'll find all the articles in there and all the different things you can do. But the problem is it's not just one thing. There's a variety of different solutions, depending on how it got stuck. And it happens all the time.

Steve: I know. I've got several machines which I've - like older machines that have just stopped being able to be updated because something, a screw loose, occurred. And it just, like, says, okay, updates have failed. And it's like, okay, you know. And I've rebooted. And sometimes if they come up - if Microsoft comes up with a service pack, that'll sort of flush everything and make it current again, and it'll kind of come back to life. And I've sort of screwed around, rebooted a few times, and it's sort of, like, oh, look, it's fixed itself. But it's, I mean, it is black magic.

Leo: Yeah. They'll get stuck. Here's an article updated on Halloween 2011, appropriately, Rev. 7.0 of this article. It's at the support.microsoft.com site, Article No. 906602: "How to Troubleshoot Common Windows Update, Microsoft Update, and

Windows Server Update Services Installation Issues." And it's really kind of a meta article that will link to a lot of other articles that will help you a little bit. And there's, as you can see, there's, well, let's see, I can't even count the number of related articles, various error codes and so forth. So there you go.

Steve: Yup.

Leo: Windows Update stalled, Windows Update failed, Windows Update blocked, all of that stuff. There's a lot of different ways to do it. Most of the time you just edit the registry to say, hey, start over.

Steve: Or there is the reformat command.

Leo: Well, it's funny because somebody in our chatroom said, "I had this happen to me, and Microsoft sent me a Windows 7 DVD and said, 'Just reinstall.'" It's a good solution.

"Gord" corrects Episode #332 about the HP calculator. He said: You said, "Ever since I was in high school, I

spent the $400 I saved up from a summer job to buy myself the HP-41, which was the very first scientific calculator HP produced. So I have long loved those machines." He says: Steve, actually it was HP-35 was the first scientific calculator. He knows, because he still has it on his desk along with his Pickett metal slide rule. Yeah. The original HP-35s were distinguished as they did not indicate the 35 on the front label. It was only after the introduction of the 45 did they have to put "35" on the old ones. I also have a 15c I still use on a daily basis. He must be an engineer. Regards - or he's a nuclear scientist. So it's the 35. 35 is the legendary HP calculator.

Steve: Yep, and that was what I purchased.

Leo: Of course it was, yeah.

Steve: In the magazines they had - it was so revolutionary that they had a actual size photo of the front of it. And of course, being the hypernerd that I was, I cut it out, mounted it on cardboard, and had it in my pocket. Oh, goodness, I was insufferable. And I remember practicing RPN on this piece of cardboard while I was saving up my money...

Leo: Waiting. Aren't you sweet.

Steve: ...from my summer job.

Leo: That is the sweetest story I ever heard in my life. He made a cardboard replica so he could practice. What are you, four? That's very cute. That's very cute. I had - but you must have had slide rules, too, those Pickett slide rules? Remember those?

Steve: Yeah.

Leo: Picketts were the aluminum; right?

Steve: Bamboo, baby.

Leo: Bamboo is the way I went. I'd go bamboo. But I had the - I had a couple. Picketts were cheaper.

Steve: Yep. The bamboo was self-lubricating. And of course, and it wouldn't expand and contract with temperature.

Leo: As the aluminum does, of course, yeah.

Steve: Exactly.

Leo: And the aluminum you had to use a pencil graphite to lubricate those. You know, my father-in-law, Jennifer's dad was a science teacher in high school. And he has - remember in high school they had the giant slide rules so they could teach you how to use a slide rule?

Steve: The demonstration slide rule, yup, with big eye-hooks so that it would hang above the blackboard.

Leo: She has one. I'm dying to get it. I would love to put it right here, hanging up right there. Wouldn't that be great?

Steve: Or maybe just kind of lean it against the wall, sort of like skis.

Leo: Yeah, because they're, like, four or five feet. They're big.

Steve: Yeah, yeah.

Leo: Really want one of those.

Steve: That's neat. So, yes, it was an HP-35. I still have it in its original box and packaging. It's in my memorabilia box somewhere. And it used an LED display, the old seven-segment display. And, yup. And, oh, and I did get confused because the 41, it was a - I also had an HP-41. That was the LCD calculator that had four little module slots on the back. And you could plug in, there was like a...

Leo: Oh, fancy.

Steve: ...barcode scanner and a printer and other things. And the batteries lasted way longer on that. I think that one used "N" cells, as I recall. And that was a beautiful calculator, too. Yeah, I've always been an HP calculator person. And at this point I'm not changing. TI will never get me. I'm an RPN...

Leo: TI was the cheap guy. But RPN, for people who say, no, that's confusing, it just makes sense. You put the operands first, and then you put the operator. That just makes sense. And it's much easier to program, if you think about it. It's stack-based; right? You push the operand, if you're doing one plus one, you push one. Then you push another one.

Steve: Enter it.

Leo: Enter it. And then it's just a stack. And then you put the operation, the plus, and it consumes the top two things on the stack to make the result, which it pushes on the stack. It's really - that's why they did RPN, because these were very primitive devices, and they couldn't make the programming too complicated. And that's how they did it.

Steve: Yeah, it's just - I just - it's perfect.

Leo: Yeah. Tom Walker, Littleton, Colorado has the Bonus Question of the Week: Thanks for the great show and thorough prep work you put into it. Yes, Steve does work so hard, probably harder than any of our other hosts. You really put hours, I don't know how many hours every week into making this show happen. I thank you from the bottom of my heart. He says: I'm curious. We get this question every darn show.

Steve: I know.

Leo: We need an FAQ. What are the three boxes of flashing lights over your left shoulder?

Steve: I figured once a year, for people who are watching the video...

Leo: Once a year.

Steve: I guess you must see it in the IRC chatroom.

Leo: The chatroom, every single show. In fact, I think there's probably an FAQ somewhere, like in the wiki or something, because they put a link in. They say, yes, here it is, here's the answer.

Steve: So, Tom, those are reproductions of the classic original Digital Equipment Corporation (DEC) PDP-8 mini computer, which I built a few years ago from a kit which was made available. And I participated in the creation of the kit, and actually a number of our listeners purchased them and built them also. So I have them running because that's pretty much all they're good for is flashing...

Leo: Is blinking lights.

Steve: Is blinking the lights.

Leo: Very fancy blinking lights.

Steve: But I grew up watching "Voyage to the Bottom of the Sea" with the Seaview and the computer and "Time Tunnel" and "Lost in Space."

Leo: Yes, me, too.

Steve: And you've got to have banks of blinking lights. That's just part of the thing you need. So I figured, in my own little way, I have some blinking lights going on.

Leo: Aren't you cute.

Steve: And every so often I look at them, and they warm my heart a little bit. It's like, aw, remember them.

Leo: And you have a whole page dedicated to these on your website, by the way.

Steve: I've got a whole subsite with the code that I wrote and the creation of those and demo videos that show how they work and what they do. So, yeah.

Leo: GRC.com. There's a menu item that says "PDP-8." Check it out.

Steve: Yup.

Leo: And that concludes our Q&A for this week, Q&A #134. We'll do another one in two weeks, and you can always go to GRC and ask a question. People will say, well, Steve never answers my question. How many questions a week do you get? I mean...

Steve: It's okay. Actually - I guess I deleted the file. I'm actually processing the history of submissions to Security Now! because I had - the testimonials that I had posted for SpinRite, I hadn't updated them since '06.

Leo: Wow.

Steve: And so people were thinking, wait a minute, nobody has given a testimonial since 2006? What's going on? So I thought, okay, I've got to fix that. We have had 47,000 questions submitted, or submissions. Of those, 6,700-some, I don't remember the exact numbers, but 6,700 mentioned, have the four-letter word "spin" in them. So, and of course many of them are not testimonials. Some are just people saying, by the way, I got a copy of SpinRite, just blah blah blah.

But so anyway, yes, the fact is, every time I check the mailbag - and I get the mail every two weeks. I pull it from the server for the purpose of preparing the Q&A. And as you say, Leo, I do spend hours. But I just can't read the 300 submissions that I get. So I read enough to get 10 to 12 good ones like we just had. And then it's like, okay. And if I had more time, I would. But I've just got to get other stuff done. You can see I didn't even shave today because the Iowa Caucus was yesterday, and I kept...

Leo: Were you up all night watching the votes?

Steve: Waiting to see what...

Leo: It was close.

Steve: Oh, goodness.

Leo: Right now, I'm just looking, right now it's...

Steve: Eight, eight.

Leo: No.

Steve: It was a difference of eight.

Leo: It's down to one.

Steve: No.

Leo: I'm looking right now. Maybe, wait a minute, maybe was that a - no, no, yeah, it's down to one. It's as close to a tie as you can get.

Steve: That's incredible.

Leo: But it's just a caucus. I don't know...

Steve: But it's just fun to watch.

Leo: Yeah. It's a straw poll. It's not in any way anything binding. But it is fascinating that it was that close. Especially since Romney spent a lot of money there, and I think Rick Santorum spent, like, nothing.

Steve: Yeah. But he spent a lot of time. Rick went all over.

Leo: He went to every county in the state of Iowa. Google may say eight, but I'm looking at - it does say, CNN - see, I don't know if we're looking at - CNN might be going, you know what, they're going back through their coverage, that's what it is. So it was one, and now it's eight.

Steve: Oh, okay.

Leo: I don't know. Amazing. What's the difference? What's the statistical difference?

Steve: Out of, what was it, it was 30 - it was a little high. They both got more, Santorum and Romney both got a little more than 30 each.

Leo: Yeah. So that's - and that's just the two of them.

Steve: Yes.

Leo: Wow. You'd need more than a slide rule to - a slide rule would not give you enough precision to calculate that.

Steve: So if anybody's interested, the next Republican debate is this coming Saturday, and then the one afterwards is the following Sunday on NBC's "Meet the Press" on Sunday morning. But ABC is doing a Republican debate next Saturday. And we're doing our next podcast next Monday. So...

Leo: Do that in Reverse Polish Notation.

Steve: Exactly. So that one's going to be on WPS, the WiFi Protected Security problem

Leo: Oh, good. Oh, good.

Steve: Which caused us last week to immediately advise everyone to disable it if they can. We're going to go in detail and in depth into how it works and how it's broken. So that's our topic for next week's Security Now!. But we are recording, as you reminded me, Leo, and we'll remind our live listeners, at 9:00 a.m....

Leo: A.M. Pacific.

Steve: Monday, Pacific time.

Leo: Yeah. That's noon Eastern, and it's 20, no, I'm sorry, nine plus eight is 1700 UTC, if you want to watch live. But, you know, you don't have to watch live because we make audio and video versions available for download after the fact. Steve's got 16Kb versions, if you want the tiniest thing. Well, no, there's something even tinier. He also has text transcriptions.

Steve: I do.

Leo: So that's the smallest version, if you want to just read Security Now!. We have the audio and video on our site, TWiT.tv. But if you go to GRC, do check out SpinRite, the world's best hard drive maintenance and recovery utility, and all the free stuff Steve gives you, including the information about Password Haystacks, ShieldsUP!, and on and on and on. Steve, thank you so much.

Steve: A great pleasure always.

Leo: Next week, WPS.

Steve: And thank you for moving the podcast forward before you take off to CES, the Consumer Electronics Show, where you guys will be broadcasting all next week for - live, I guess, right?

Leo: Yeah, tons, all day and then part of the evening, too.

Steve: Cool.

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

Steve: Thanks, Leo.