Listener Feedback #239

Description: Leo and I discuss a bit of Flip Feng Shui follow-up; Apple's announcements; Android's rough week; wireless device privacy leakages; some fun miscellany; and 10 questions, comments, and observations from our terrific listeners.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-577.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-577-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve's got the latest security news for you. And, finally, it's a Q&A session. So we've got a lot of questions from our audience. Steve has a lot of answers for you. Coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 577, recorded Tuesday, September 13th, 2016: Your questions, Steve's answers, #239.

It's time for Security Now!. I'm going to get my Temperfect mug and put in my Bragi headphones and listen to this guy right here, Mr. Explainer in Chief himself, Steven Gibson. Hi, Steve.

Steve Gibson: Say, you know, I have had some people ask why I'm drinking, you know, those who see the video see me holding up a white ceramic cup.

Leo: Yeah, a Lavazza cup.

Steve: As opposed to the - yeah, Lavazza, as opposed to the Temperfect mug. And it's just that, if I were out and about, I would probably use the Temperfect. But for my mode of operation, it's just easier for me.

Leo: Folks, no cup of coffee lasts more than a few minutes with Steve Gibson around. There's no issue of it cooling or chilling. We have a...

Steve: Yeah. So there's no need for it to - and the cup is also sitting on a hotplate, you know, a traditional old school cup warmer. So it keeps the cup warm.

Leo: Oh, well, who needs a Temperfect mug if you've got a cup warmer?

Steve: Exactly.

Leo: You keep it on a hotplate, really?

Steve: Yeah.

Leo: Wow. You don't like cold coffee.

Steve: Don't think I can show it to the camera. Yeah, there. A little hotplate.

Leo: Oh. Is it one of the USB-powered hotplates?

Steve: Oh, no, no.

Leo: That's the real deal.

Steve: That whole idea, there's something that should not be USB powered.

Leo: 110-volt hotplate.

Steve: Oh, god.

Leo: Yeah, those things...

Steve: But there's like a USB fan, where you stick it into your laptop and then...

Leo: Those are all right because you don't need a whole lot of wattage for a fan.

Steve: That doesn't need a lot of power.

Leo: No, a hotplate, yeah, maybe not.

Steve: Well, so - did we start?

Leo: Yes. Yeah, we started, yeah, yup.

Steve: This is Security Now! 577. Finally we have a Q&A.

Leo: Woohoo.

Steve: The industry has been sufficiently quiet, and the pressure to do a Q&A has been growing. I have, Leo, I have so many questions that we could just, like, we could cruise for another couple of decades just doing Q&As because so many people write. And between only really having time to do 10 and giving them useful coverage, plus whatever news has happened in the week, and the fact that sometimes we just can't get to a Q&A because of other really cool stuff that we have to talk about, I have a massive backlog.

Leo: Well, the truth is, every show you do stimulates more questions than you can answer.

Steve: Yeah.

Leo: I mean, I know myself, as I'm listening, I'm thinking of questions. But, okay. So you - and then what - and so I can understand why there's...

Steve: So we're going to - I have a quick little fun Flip Feng Shui follow-up from the authors. Apple's announcement I wanted to talk to you about just briefly. Android had a rough week. There's also some more wireless device privacy leakage in the news. We have some fun miscellaneous stuff; and, of course, being a Q&A, 10 questions, comments, and observations from our terrific listeners. So a great show.

Leo: Excellent. And a great Image of the Week, which we'll do in a moment. I love it.

Steve: So our Picture of the Week is just fun.

Leo: I love it.

Steve: A lot of people have had fun over the years spoofing various O'Reilly covers. I don't know why...

Leo: The animal covers, yeah.

Steve: Yes. Yeah, they're just notorious within the IT, computer, and programming community. They're generally great references. I have, you know, not a complete set.

You can't have a complete set. There are just too many.

Leo: I have a ton of them. I'm looking at a bunch right now. And the animals have become so much their trademark that many of their books are known by the animal on it.

Steve: Yeah, what's the Perl?

Leo: So the Perl book is the Camel Book.

Steve: That's right.

Leo: Yeah, Larry Wall's Perl book is the Camel Book.

Steve: Yup.

Leo: That's how successful they are. I've never seen a giraffe book.

Steve: Well, yes. And so this was well done. Anyway, it's just fun. It's a fun cover: "Regex by Trial and Error."

Leo: Which is how I do it. I don't know about you pros, but...

Steve: But I think, no, I retweeted - I guess I shared it with a couple friends who are also programmers because I think everyone does. And that's the point of this, is that the regular expression language is formally defined. There unfortunately is not a single one. All the different implementations, authors couldn't resist changing this or that. So, for example, the Regex bible talks about the individual, has individual chapters for different languages' implementations of regular expressions, Perl of course being probably the preeminent one. But it's really - what's the term? It's not procedural, it's declarative. And I would hate to have to code the interpreter for regular expressions. I have sometimes marveled at what the regular expression interpreter is able to do.

But the idea is you're able to describe the way you want typically a string to be transformed with very complex pattern matching and substitution phrases in this crazy thing. And it often doesn't work right the first time. So you do it. Then you give it some test cases. And then you say, oh, whoops. And in fact I'm sure the world is full of bugs that have not yet been found in regex statements which were not quite what the programmer intended for them to do. So anyway, just - so "Regex by Trial and Error."

Leo: By O RLY books. I love that; right?

Steve: There are a couple of typos on the cover that were no doubt a consequence of a

misapplied regular expression.

Leo: "Combining slashes and dots until a thing happens." Love it.

Steve: So I was surprised and really pleased. I received a tweet from one of the security research team at VU Amsterdam's group that's run by Professor Herbert Bos, who we talked about. They're the guys who did the Flip Feng Shui work which was so impressive. Ben Gras wrote, he said: "As one of the authors" - he tweeted to me - "one of the authors, thank you for your knowledgeable and detailed exposition. We're honored by it and your kind words." And I replied, I said: "Oh, wow, thanks." And so in email he followed up, saying: "Thank you again for your Flip Feng Shui coverage. Because of more recent coverage, I am a bit of an expert on how well our work is understood; and I don't hesitate to rate your coverage at a 95th percentile rating for expertise and quality of exposition."

Leo: Wow, that's excellent, wow.

Steve: "We also loved your appreciative words, of course. So full marks, thank you so much." And then he concluded with: "I'm also a big fan of the show, so it's just joy all around here."

Leo: Oh, that's great. High praise, yeah.

Steve: So neat to know that these guys are following the podcast. And thanks, Ben, for the note.

Leo: You get used to, I'm sure, if you're in a project like that, of kind of incorrect or, I mean, mainstream media just - mainstream. The tech media just mangles stuff.

Steve: Well, how often do I talk on the podcast about something that's just like, okay, this is not a problem; or, oh, this is much bigger than expected. And typically the press tends to overheat these things, as we've seen, because...

Leo: It's good for ratings.

Steve: ...they're looking for clickbait, yeah.

Leo: Yeah.

Steve: So I did want to just mention briefly about the Apple announcement last week, last Wednesday. I slept soundly through Thursday night into Friday morning for the first time in quite a while, not bothering to set an alarm.

Leo: You did not get up at midnight?

Steve: No, honey.

Leo: I did.

Steve: Now, and I also...

Leo: I did.

Steve: I know you did.

Leo: I had to.

Steve: And I guess I would have been frustrated because, if I was going to get up at midnight, I would have wanted to get the product I wanted.

Leo: Yeah, no.

Steve: And you did, and you were unable to get what you wanted.

Leo: Not until next - not until I get back from vacation, let's put it that way.

Steve: No, but I didn't think you were able even to order the phone that you wanted.

Leo: No, I did. The reason I ordered two, I ordered a jet black, which will get here when I get home from vacation.

Steve: I see.

Leo: October 4th.

Steve: And then one to take with you.

Leo: But I wanted, well, not just to take with me, just so I could review it this weekend on the radio show.

Steve: And this is for the trip where you're not taking any technology along.

Leo: Yeah, by the way, that didn't last.

Steve: Okay, good.

Leo: Now I'm bringing more technology than ever.

Steve: So I just wanted to observe. I mean, you've had some great things to say about the nature of what happens as companies and products age. And this makes me think of Windows. And that is, I'm staying with my 6 because, you know, I have a 6s Plus. I followed them until I got everything I wanted. I now have everything I want. Just like, yes, a faster processor would be nice. A better camera would be nice. More memory might be nice, although I've got the biggest one at the time, which I haven't had a problem with. But I don't need that stuff.

And I did sort of think it was interesting last year when Apple said, oh, sign up for the new phone every year plan. And it's like, uh, no. So maybe I'll go to the 7s. Or maybe I'll wait to see what the 8 is. My point is that this is what happens at some point when things are good enough. And Windows has been good enough for many generations now, which is why many people are choosing to stay on 7.

Leo: Right.

Steve: That is, Windows 7. And it's very difficult, then, for companies to move people forward. And so, in my opinion, the iPhone has achieved that same status. And that is, it's like Word years ago was done. But "done" doesn't make them continuing profit. So they had to keep cranking out new versions and try to upsell us. Except, sorry, this just does everything I want. And I understand that. So I just think that Apple's at the same place. And it'll be fun to see how the whole lack of a headphone jack thing shakes out. I did hear some corroboration of the argument that space inside the case is at such a premium that the amount of space required for the interior management of that little hole, that little penetration through the case is surprising. And that the designers of the mechanics would love to have that back. And I did note one of your guests noted that Bluetooth consumed energy, but failed to note that so do headphones. And so that's probably a wash.

Leo: Oh, yeah. We don't really think about that. But of course you've got to drive those speakers.

Steve: Correct.

Leo: Think it's equal? That's an equal amount?

Steve: And in fact I believe that the battery voltage is so low that it needs to be stepped

up and provide enough pressure to drive the impedance of the headphones.

Leo: Oh, that's interesting.

Steve: So that's lossy also. So the headphones actually are a problem. And you get so there are a lot of benefits associated with just saying, eh, we'd rather just use radio because we already have radio. It's already there. It already works. We're just going to send the audio out that way. I mean, I'm not discounting the controversial nature of it. I'm glad I still have my headphone jack. But I just use - I haven't had a conversation on my phone for years.

Leo: I know.

Steve: I shave much more often than I talk on the phone.

Leo: But, see, you probably don't listen to audio in your headphones, I mean your phone, that much.

Steve: No. I just - I don't [crosstalk].

Leo: Yeah. So many, many, many, many people, it's their iPod. I listen to audiobooks, as well.

Steve: Yes, yes.

Leo: I listen on - and I have - and also to the point where I've purchased much higher quality headphones for my phone. So it makes a big difference to me.

Steve: Traveling, if I'm flying...

Leo: Same thing, yeah.

Steve: I use Bose sound-canceling headphones and my phone to...

Leo: Can you use Bluetooth on a plane? No, that's a radio.

Steve: Oh, yeah. No, you're right, I plug it in.

Leo: So that's another argument against the thing.

Steve: Of course, those arbitrary restrictions are getting relaxed because people are saying, you know, we're not going to fly on your plane unless we can use our stuff.

Leo: Right.

Steve: And the airlines go, oh, gee, I guess it's not such a security problem after all. So speaking of security problems, Android has had a rough week. Google just released patches for their Nexus devices for two serious problems, the worst of which was a - and we've talked about this class of problem before - a parsing problem in JPEG images, in this case in the - do you pronounce it EXIF?

Leo: Yeah, that's right, E-X-I-F.

Steve: EXIF, E-X-I-F. That's the metadata for the JPEG. And we've talked about how any image is sort of a compiled thing, and you need an interpreter in order to display it. And so, and that code is tricky to get right. There's probably regular expressions in there somewhere. And so vulnerabilities were found. And then what's worse is that applications were discovered in the Google Play Store, downloaded as many as 2.5 times - 2.5 times - 2.5 million times.

Leo: Yeah, that was a shocker, wasn't it? Wow.

Steve: Yeah. So it's like a whole bunch of people have apps that have this malicious stuff in them already. So Google's fixed it. And of course this is the problem with this really long tale of upgrading on Android is that Google has made the patches available to their OEM partners, but no one knows how long it's going to take for those to get pushed out. We watched, because everyone could see with Stagefright, how long it took the various manufacturers, even though Stagefright was a code red level emergency. We had tools that allowed us to check our Android devices, and it took weeks and in some cases months for Stagefright to get fixed, when it was fixed at all.

And so of course the bigger problem is many people are using older Android devices that just won't ever get patched. It's a computer. And as we've argued on the podcast, there has to be some sort of responsibility that goes along with offering connectivity services to a computer like that. As long as you're profiting from its connectivity, I would argue you have an obligation to provide patches when they're created. But the Android ecosystem is in the process, I think, of improving. But it's had a rough time.

There was an interesting article, I think you might have talked about it on TWiT on Sunday, about the Bluetooth Low Energy leaking information. Or maybe it was you were talking about WiFi. But actually this sort of pulls it all together. There was an interesting article. And I did have to chuckle because of a little bit of an anecdote in here. Sean Gallagher, writing for Ars Technica, I'll just read - I'll share the beginning of his story. He wrote: "My new neighbor was using AirDrop to move some files from his phone to his iMac. I hadn't introduced myself yet, but I already knew his name. Meanwhile, someone with a Pebble watch was walking past, and someone named 'Johnny B' was idling at the stoplight at the corner in their Volkswagen Beetle, following directions from their Garmin Nuvi. Another person was using an Apple Pencil and their iPad at a nearby shop. And someone just turned on their Samsung smart television.

"I knew all this because each person advertised their presence wirelessly, either over 'classic' Bluetooth or the newer Bluetooth Low Energy (BTLE) protocol. And I," writes Sean, "was running an open source tool called Blue Hydra, a project from the team at Pwnie Express," P-W-N-I-E. "Blue Hydra is intended to give security professionals a way of tracking the presence of traditional Bluetooth, BTLE devices, and BTLE iBeacon proximity sensors. But it can also be connected to other tools to provide alerts on the presence of particular devices.

"Despite their 'Low Energy' moniker, BTLE devices," he's writing, "are constantly polling the world even while in sleep mode. And while they use randomized media access control (MAC addresses), they advertise other data that is unique to each device, including a universally unique identifier," the so-called UUID. "As a result, if you can tie a specific UUID to a device by other means, you can track the device and its owner. By using the Received Signal Strength Indication" - that's RSSI - "you can get a sense of how far away they are. That information," he writes, "can be used for good or ill, to generate movement data about the people who carry those devices, and to watch for devices that appear when they shouldn't. Pwnie's Rick Farina told Ars, as he gave us a walkthrough of the tool: 'I have an alert set up for when my mother-in-law's car pulls into range. It gives me about a 30-second warning.'" So it does have some beneficial applications.

Anyway, so I just wanted to note that what we're seeing, this is another instance of the tradeoff between absolute privacy and features, which is a problem with all of our wireless technologies. As we've covered in the past, WiFi similarly blabs about who it is, where it's been, and who it knows. And we very much want the convenience of having everything just work. And I've noted that, when things just pair instantly and easily, everyone cheers. Oh, look how easy that was. And it's like, yeah, but what we forget is that the nature of the underlying interchange which participates in making that so easy in many cases is doing things like using beacons to broadcast their presence and, often inadvertently, the identity of their owner.

So I don't think I see this trend reversing. It seems that we're seeing more of this, even in a privacy- and security-conscious environment now. But people want the convenience. And look at the explosion. A perfect example is of IoT devices and the coverage we've been giving this year with the appalling lack of security because features trump security, for a long time. And we're learning lessons first from the desktop and now from smartphones about how important it is to maybe rein in features, or at least really focus on security because, if you don't, you're just not going to have any.

Many people tweeted me this story because, of course, my interest in hard drives. You may have seen this story, Leo, about the datacenter at ING, or maybe it's just NG Banks. I think it was ING. I just wrote NG in my notes.

Leo: I want to say ING, but I can't remember.

Steve: Yeah, I think it's ING, yeah. Their main datacenter in Bucharest, which is in Romania, was down for 10 hours, brought down by a seemingly unrelated test of their high-pressure fire suppression gas discharge system. So everyone, everything was fine. No fires. But they said, "Let's test to make sure that the fire suppression system works." So they opened the valves. And the sound made as the gas was pushed through huge numbers of tiny holes, first of all, it pinned the meters, the sound level meters, which only went to 130dB. It pinned them. So we don't know how loud it was, but it was more than 130dB.

And as we've discussed before, hard drives have become so sensitive to mechanical vibration. Remember years ago we showed the funny YouTube video of some guy shouting at his hard drive RAID array and making it go wonky. I think we were seeing a graph of the data transfer rate. And when he screamed at it really loudly, the transfer rate dropped because what was happening was he had, essentially, the tracks were so close together, IBM researchers say one one-millionth of an inch offset will now, in contemporary drives, will cause the drive to be unable to read its data. One one-millionth. What's that, a micron?

Leo: Well, that's a millionth of a meter. So, but yeah.

Steve: That's right. Right. So way more than that. One-millionth of a inch.

Leo: Right.

Steve: So the sound pressure of this fire suppression gas released caused these drives to go off track when they were writing. And that's the critical thing. That's what did the damage because, if you're reading from the drive, what typically happens is that the drive can't find its sector, or it gets a read error which is uncorrectable because the head's too far off track. So it goes around again and just does a retry and may be able to succeed. But if you've got a busy datacenter where you are recording information in real time, and you force those heads off track, then you're in trouble because you catch the head while it's writing. And it is then not writing in the right place. And so you can't read it again, and it may well have overwritten its adjacent neighboring track. So that's a big problem.

I think I've mentioned a couple times probably through the years that the very first version of SpinRite, it's like pre-SpinRite, because SpinRite was really born to adjust sector interleave. But doing it correctly meant I had to also make it do really robust data recovery because I was going to low level format the track and forever lose the opportunity of recovering data. So it absolutely had to be able to perform data recovery as part of the task of basically doing a true low level reformat of the drive. But the predecessor to it was a utility that I wrote, just a one-off, to fix a girlfriend's hard drive.

She had, I think it was a 10MB Seagate that all of her company's financial data was on, that it could no longer read. And it turns out that just lifting one end of the drive up, it was a voice coil-actuated drive. And just lifting one end up allowed gravity to bias where the heads settled so that they were back over the data tracks, and it would then read. And so what I did was I wrote a simpleminded version of SpinRite that essentially read and rewrote the data. And I ran it several times while gradually returning the drive to horizontal and essentially migrated the actual physical tracks back to where they should have been.

Leo: Oh, you crack me up.

Steve: Back in alignment. It worked.

Leo: Oh, my goodness. Wow.

Steve: Yeah, it worked.

Leo: That's crazy.

Steve: So anyway, another friend of mine had noted that he'd had problems with servers that were seeing too much vibration. And they found, that just putting their thumb on the server damped the vibration enough that things worked better. So this is something that everyone's going to have to pay some attention to, those who continue to use spinning hard drives, is that, with this crazy density comes some responsibility. The drives really have to be treated carefully, especially when they're writing, because that's what had this datacenter down for 10 hours. They had to recover because that sound did lasting damage to their drives.

Leo: Amazing, yeah.

Steve: Yeah. And also in a little bit of miscellany - I have two more. I just thought to note, and you guys have been talking about it on many podcasts, you know you're in trouble when the headlines read: "How to tell an explosive Galaxy Note 7 from a non-explosive one."

Leo: How, pray tell?

Steve: Yes. Or worse, when the headline reads: "Samsung Galaxy Note 7 explodes in New York, burns a six-year-old boy."

Leo: Wow.

Steve: So, and actually later in our Q&A we have a little discussion about - because there's some confusion about proper handling of lithium-ion batteries. It's something that we've talked about here several times. But I agree with you, Leo, that Samsung has done everything they can, which is to essentially recall a huge number of those very nice phones and take responsibility for them.

Leo: Yeah, yeah.

Steve: And the fact that these things are able to do that reminds us that there's a lot of energy, a lot of chemical energy stored in those cells. I mean, the best minds in the energy storage business, in the battery design business, their whole goal has been, for many years, how can we possibly cram as much energy as possible into the smallest size with the lowest weight? And what we have today in lithium polymer, lithium-ion batteries, is the result of our best engineering, storing energy in that medium. And

unfortunately, when it comes out all at once in any form, that's not good. So we want to bleed it back out gently and use it, rather than in any way have it suddenly present itself all at once, as happens when the battery structure breaks down and basically just it uses itself, it uses that energy against itself and raises the temperature. Gas is released; you get an explosion, and a big recall of, I mean, a big, expensive recall.

And one last thing. At TechCrunch, I think it was TechCrunch Disrupt, there was a super high-performance database that was announced. And I loved that they used GPUs to accelerate their SQL-compatible database. All we ever see is GPUs used for cracking things, pretty much. They're, like, they make massive hashing engines. And of course before custom silicon was available for our various crypto currencies, GPUs were used in order to perform hashing. And then those got replaced by custom silicon.

But this is called the BlazingDB. If anyone's interested, it's at BlazingDB.com. And I also got a kick out of the fact, first of all, they tout the fact that it's between five and 140 times faster working on enterprise-class and enterprise-scale databases. And then they boast that it's written in C/C++, calling it "a low level language with very granular control of hardware - memory, processors, et cetera." And then they say: "C/C++ is our dedication to delivering massive scale and hyper speed for our customers."

And I sort of thought, okay. I guess that's probably, when you compare C, as we've talked about over the last couple weeks, to various higher level languages which represent much greater levels of abstraction from the low level, you know, C was designed to be an OS implementation language, just far enough away from the actual hardware of the machine to give you machine independence, yet without much cost in performance. So anyway, I thought it was cool to see our ubiquitous GPUs being put to a nice non-crypto and non-graphics accelerating task.

And speaking of C, I needed to correct the record. I said incorrectly last week that Doom was written in assembler. Turns out only a tiny bit was written in assembler. Most of it is written in C. I got a note from James Boer in Kirkland, Washington, under the subject "Doom Trivia." And he wrote: "Just FYI, Doom was written in C, with some small portions of rendering code in assembly." Which is exactly right. I went and looked.

He said: "Even so, you're correct that it was masterfully coded. Doom was partially what convinced me," he writes, "to become a professional videogame programmer. Doom was not the first pseudo-3D rendered game, nor was it the first shareware game. What made it unique, among a few other details, was that it was one of the first games that could render reasonably complex world geometry using the entire screen on a 486-class machine - and, of course, for the stunningly visceral game play. In case you're wondering why I call it 'pseudo-3D,' that's because it's not true 3D rendering like Quake" - which of course came after - "and is sometimes referred to as '2.5D.'"

And then I didn't realize, but GitHub has id Software's source. And I have the link in the show notes, if anyone is interested. It's at github.com/idSoftware/DOOM. John Carmack posted the entire source of Doom up there. And in fact, as is often the case, there is a chunk of assembly in the most time-critical portion because nothing beats hand-designing code for speed. But I was very impressed that as much of it was in C as is. I mean, virtually all of it is in C, with just like a page and a half of assembly code down in the actual texture mapping and rendering portion of the code. So, James, thank you for the correction. I'm glad for that.

Leo: This is great, to have the source code.

Steve: Isn't that neat, to have the source of Doom?

Leo: Yeah.

Steve: Yeah, I mean, anyone who was a budding game programmer, there are, represented in that code, a huge number of really great ideas. And I was curious. I have a - Michael Abrash was one of the great low level coders of yesteryear. And I checked because I was curious. I have his - he did the "Zen of Assembly Language Programming." I have all of his original books, and also the "Graphics Programming Black Book Special Edition," with a foreword written by John Carmack, where John talks about having tried many times without success to lure Michael away from whatever he was doing during the time that Doom and then Quake were being put together because these guys are masters of the low level arts. Great coders.

Leo: Yeah. Yeah. Really neat.

Steve: And I have a nice note from Paul in Worthington, U.K. I just loved how he put it together. He shares his recent SpinRite success with SSDs, which he wasn't aware of. He said: "Steve, I cannot remember," he writes, "how I came across SpinRite. Think it may have been back at v5 then. I assumed when I switched my system drives to SSD that I would be using SpinRite a lot less. I used to have the habit of running SpinRite regularly on my system drives every few months, or if the system started to slow. It always was great, and I have recovered drives that blue screened or failed to start completely.

"I built a new machine last year using a SanDisk SSD. I chose an enterprise version with a 10-year guarantee" - this is last year - "because, as a commercial photographer, I need speed, but value reliability even above speed. It was great. SanDisk dedicated software allowed me to monitor it and run a TRIM command. I had a big project running," he says in parens, "(layered Photoshop files sized between 3 and 7GB on disk), and the system slowed, particularly on startup. But some days it was fine, fast and no problems. SanDisk emailed to tell me I should update the firmware, but backup everything on the disk first, and also update the dedicated SanDisk SSD Dashboard software.

"First, I updated the SanDisk SSD Dashboard software, and it FAILED COMPLETELY," he has in all caps. "It just destroyed the existing working version. No worries, just roll back the system." And he says, "I use FarStone recovery software. No good. Install the old version of SSD Dashboard, another no. Okay, I should have downloaded and copied it like I always used to, but these days you don't need to, do you?" he writes. By this time I'm tearing what is left of my hair out. So of course I start thinking, 'What I need is SpinRite for SSD.'

"To cut a long story short, I dug around the Internet and discovered that SpinRite can help SSDs. OF COURSE IT DID," he writes in all caps. "It found one defective and unrecoverable sector close to the beginning of the drive, and two further ones that it recovered. Rebooted the machine, and since then it's been running like lightning again. I, for one, know that SpinRite is one of the most cost-effective software purchases I have ever made. I look forward to its future fixing both my spinning and my non-spinning mass storage. Thanks for your great software and the Security Now! show." And, Paul, thanks for sharing your experience.

Leo: Awesome, awesome. All right. You want to do some questions?

Steve: Let's do it.

Leo: I've got them right in front of my little eyes right here. Let's start with JK in LV. Las Vegas? Probably.

Steve: That's what I was wondering.

Leo: I'm thinking. Could be Little Village. I don't know. Wants to know about the need to wipe an encrypted drive. Is there ever a reason to do a secure delete on a hard drive that's been using whole disk encryption? I think it's dumb, no matter if an SSD or spinning drive. My friend says, "Oh, no, you're wrong." I'd be happy to be told otherwise, if the Security Yoda disagrees. Thanks for everything you do. Love Security Now! and SpinRite. All right, Security Yoda. [Yoda voice] What drive must we delete?

Steve: So there's two different ways today that we have whole drive encryption. One is if it's built into the hardware of the drive, and the other is if we add whole drive encryption afterwards, that is, on top of an unencrypted drive. And I think the second, from reading the question, I think the second is what he was referring to. And the advantage of the second is that we know exactly how the system works. Where the encryption is built into the drive hardware, you just can't know. And everybody who listens to the podcast knows how I am about details. That's where the devil is. It matters how they implemented it, whether it is truly secure when you change or lose the password or delete it or do a secure wipe or something.

The advantage of adding our own is that we know how it works. And it is absolutely true. So JK is right. His skeptical friend, I would argue, is incorrect in that, if you have encrypted, you've added your own whole drive encryption to a hard drive before you have started using it. That's kind of important because, if you're really super concerned, because if bad sectors were spared out, that is, removed from service, and while - I'm sorry. If they were removed from service while the drive was unencrypted, they would forever remain inaccessible to the encryption and unencrypted. There's not going to be much data there. There's going to be 512 bytes, typically, of data because we typically spare out on a sector granularity basis. But still it's not encrypted.

So if you know you want security, encrypt the drive immediately, then start putting your data in. If you then delete the key, maybe overwrite the header, I don't remember now if TrueCrypt has an explicit "expunge the header." Because that's - it's in the header, which actually is stored redundantly because it's so important. You want to absolutely securely delete that because that's where the master key is which your password is used to decrypt, which then allows the drive to access itself, essentially. But with that gone, and proper encryption, it's just pseudorandom noise. We understand how to do that now. We've got that technology nailed. Without the key and contemporary encryption, there's no way to reverse that data.

Leo: You don't have to worry about the swap drive or anything like that?

Steve: Well, all of the good technologies now do that.

Leo: They encrypt the swap, as well.

Steve: They will encrypt the swap drive and encrypt the whole boot process and everything. The hibernation file and the swap file, also. So that said, the data is still there. It's encrypted. It's inaccessible. If you really are a belt-and-suspenders kind of person, yeah, run DBAN, Darik's Boot and Nuke, over the drive. But don't do it 15 times. Just write zeroes. That'll be plenty.

Leo: [Yoda voice] Encrypt or do not encrypt. There is no [indiscernible]. That was terrible. I'm sorry. I apologize. Brian in New Haven, Connecticut has our next question, wonders: How can I have told the difference between the Ubiquiti EdgeRouter X and a managed switch?

I have appreciated your discussion of the Ubiquiti EdgeRouter X, our little $60 miracle worker, over the last several episodes of Security Now!. I can understand now how that device provides five separate logical interfaces that allow for network isolation in a way that the other blue box routers containing a simple switch cannot.

But if I had just been browsing Amazon, and the EdgeRouter X appeared among several managed switches, I doubt I could have discerned the difference on my own. Is there anything in the listing that could have made it more evident? The mention of 5GB RJ45 ports? Seeing the ports labeled "eth0" through "eth4"? I'm trying to figure out if there was something obvious I missed, or if the Amazon listing obscures this capability?

Thanks for the excellent podcast. Been a listener for at least seven years and owned my copy of SpinRite for nearly that long. Wish I had an anecdote to submit, but maybe I don't because I run it regularly on all my drives.

Steve: And you know, I forgot that I should have noted - the guy that's provided the nice testimonial about SpinRite on his SSD, if you're listening, do run SpinRite, now that you know it works, on that SSD. Just use Level 2, which is the read-only pass. And that'll still help the drive to keep itself in good shape. So definitely worth doing from time to time.

Brian's question was a good one, and it made me think of the time I spent digging into - what was that other, it's not MikroTik.

Leo: You liked the MikroTik. That was a good solution.

Steve: MikroTik? I don't remember now what the right way to say that was.

Leo: No, I think it's like MikroTik.

Steve: Anyway, it just - it was so unclear from all the descriptions - they had, like, 75 different routers in the first place. It's like, okay, really, do we need this many? But it just wasn't clear. So I had to dig into the chip in order to see what these things actually did. And that's when I discovered that the chips that even incapable routers were using were capable of doing this. But they just hadn't bothered to. So the answer, Brian, is, and to our listeners, no, it's not obvious. And it would be nice if manufacturers understood that this is a feature we care about. I mean, they talk about things like power over Ethernet and that MDIX where it doesn't matter, you don't have to have crossover cables anymore. The jack is able to flip the connections around when you're jumping between ports that have the same sense, whereas we used to have crossover cables. All that's kind of gone away. They talk about those things.

But the problem is there are managed switches which give you, for example, filtering on ports, but not the ability to define disjoint subnets. And they're not routers because, if you have multiple networks, that's a classic definition of a router. So maybe the distinction, although I don't know that you could hold them to this, is that the Ubiquiti EdgeRouter, the key is the word "router." A managed switch might have that capability. For example, that Cisco, that SG300 series, it's technically called a "managed switch." It's not called a "router," but it does allow you to set up different subnets on that box.

So, unfortunately, Brian's identified a gray area. And the only thing I know you could do would be not to look at the bullet points, unfortunately, because they're just not going to be clear. Go track down the manufacturer's site. Maybe they say more. Maybe look through any social media postings that are hung onto reviews to see what other people are saying, whether they were able to use that or get that working. And check out the manual. Grab the PDF of the manual and see if it gives you that feature. Unfortunately, there doesn't seem to be an easy way to do this. We do know that the Ubiquiti EdgeRouter X does the job. And I just haven't found a reason, unless you needed more than five ports, for choosing a different piece of hardware. At 50 bucks, it's a real bargain.

Leo: Some concern has been expressed about Ubiquiti's privacy policies and so forth. I haven't really - I haven't looked into it.

Steve: I saw that go by. And I don't remember why I dismissed it now. But I sort of thought, oh, okay, that's interesting.

Leo: The problem is all terms of service will have these overly broad, for legal reasons, overly broad things like we could use this information in a variety of ways. And they're just protecting themselves against [crosstalk].

Steve: And often it's the corporate attorney who insists on putting this language in there. I remember, I've always had this policy that someone could use SpinRite on all the drives they owned. But my printed manual did not say that. And I argued with my marketing and sales department, back when I had them. Now I don't, so the problem went away.

Leo: End of argument. End of argument.

Steve: It's like, no, I'm not going to ask someone, I mean, you're crazy to think anyone is going to buy it four times if they have four drives. That's nuts.

Leo: Right, right.

Steve: So I won.

Leo: Commonsense is uncommon.

Steve: Yeah.

Leo: Especially when you get a lawyer involved. Aaron, who is @vader in real life, @vaderIRL on the Twitter, needs to - [Yoda voice] Darth Vader he is - need to use a VPN and wonders whether Steve still recommends proXPN. Actually, I don't think you ever recommended proXPN. That was an advertiser, so let's make that clear, that advertising endorsements are not the same. They come from us, from me, and I don't ever expect Steve to get involved in that.

Steve: Well, and I sort of have an interesting take because things have changed. In our contemporary modern surveillance world, the traditional centralized VPN server model, I would argue, has become maybe a little challenged. Certainly valuable, but I would say there's certainly a use case for it. But the problem is, if what people want is true privacy and surveillance avoidance, the concern is that this is not unlike the trouble that Tor exit nodes are known to have.

We know that intelligence and law enforcement agencies are naturally attracted, sort of like bees to honey, to Tor exit nodes because that's where the information is. Something is coming out of there that somebody wanted to obscure somehow. They went to some effort in order to hide themselves, which sort of begs the question, "Huh, wonder what's going on there?" And a VPN's exit node is essentially, it is a data concentrator by nature. It's inherently similar. There are people using a VPN for whatever reason. But at that server point, the traffic emerges unencrypted from the VPN encryption tunnel that was carrying it, out onto the Internet, where it is then subject to scrutiny. So as opposed to the pre-Snowden world, today's hugely increased, nearly de facto application of TLS encryption for all web communications does, I would say, dramatically reduce the need for a separate encrypted tunnel in many instances. Not all.

And the problem is we can't know that all of our traffic is encrypted unless we explicitly take responsibility. But so, for example, the typical model of operating at open WiFi at Starbucks, or a hotel's network, which we used in the past as examples of horrifying problems, where you just - it's amazing to see the amount of plaintext going by. That level of plaintext has dropped to almost nothing because encryption is becoming, post-Snowden, there's been a huge move in that direction.

And so for many applications I think that running one's own VPN server at home can

make much more sense. Then, when you're out on the road, your traffic can be protected on its way to your home base, where you are then able to directly access your home assets - like, Leo, your Drobo, which you've left at home with your 300 Audible books, and you can grab one that you forgot to download. But also your traffic can emerge onto the Internet from there, even if you're traveling remotely.

So what that avoids is the attention concentration that any commercial service creates. And of course, as we know, home routers are increasingly supporting OpenVPN natively. The pfSense firewall offers it because it's FreeBSD based, and OpenVPN is a feature, a dropdown menu feature of it. And even though it's not available at the UI, we know that at the command line even the Ubiquiti EdgeRouter is able to run OpenVPN. So there's, you know, not only is it providing you with really good security, it's also giving an Internet-facing OpenVPN service that the user can access wherever they are in order to get their traffic out of the environment where they're located securely to their home base. And there they have access both to their internal network in a secure fashion, as well as the rest of the Internet, without it coming under the concentrated scrutiny of any single high-volume exit point onto the Internet.

And as we're discussed before, that little controversial super simple shell script which was created by one of our listeners, that PiVPN project, for $35 you get a Raspberry Pi. Which, by the way, just passed 10 million units sold. Which is amazing when you consider it's not even a phone, it's a circuit board. Yet 10 million of them have sold. They just crossed that benchmark. So that little Raspberry Pi can be used to create an OpenVPN endpoint that will do the same thing plugged into a spare port inside of anyone's network, into the router there.

So it is the case, I think, that a VPN provider can be needed, for example, if you want an international presence. And we know that we want one that does not log our traffic. And we know that proXPN is not a traffic logger. So I do support them from the standpoint of knowing of no disqualifying factor. I think they're a good company. Although, again, I'm beginning to wonder whether that traditional traffic concentration model holds up as well as an individual personal VPN which allows you to access the assets you have at home or the Internet, no matter where you are.

Leo: I might disagree just slightly with you. It depends on what your goals are.

Steve: Correct.

Leo: Not everybody is trying to avoid government surveillance. I mean, that is just one of many reasons to use a VPN. So, yes, if you don't want - if you're trying to avoid government surveillance using Tor or VPN, or PGP for that matter, is a red flag. And so you might be attracting attention. But as you point out, people use VPNs for a lot of reasons. You're not going to use your home VPN if you want to avoid geographic restrictions.

Steve: Correct.

Leo: You also, I think, should use a VPN in an open WiFi access point. I've said, I agreed, for a long time I said exactly what you said. Well, as long as you're using

encrypted services, it's not a big issue. But increasingly these very widely available tools like the WiFi Pineapple that Hak5 sells, which they use in other ways to attack you, even if you're using SSL websites. I mean, if you're sitting in an open WiFi access point, they can observe some things about you. They can apparently, and I'm not an expert on this, even figure out what WiFi access points that your system has attached to in the past and then spoof it.

Steve: Yeah. Unfortunately, a VPN won't protect you.

Leo: Ah, yeah, because your WiFi would then be promiscuous and join this other guy's thing. You mean a VPN won't protect you against a Pineapple?

Steve: No. No, because essentially the VPN is data traffic, but you still have your WiFi presence.

Leo: Oh, I see what you're saying.

Steve: It's still there.

Leo: That's why I use a hardware firewall, the Tiny Hardware Firewall. So at that point I'm using the firewall to choose an open access point. It joins the access point. I join the firewall. The firewall is my access point on all my devices. And it also then logs me through a VPN, and optionally a Tor server, as well.

Steve: Yes. So it becomes your point of presence.

Leo: Right. And it's too dumb to be useful to anybody with a Pineapple. I think. I hope. Tell me if I'm wrong. So I think there are arguments for using a VPN that might supersede the argument, well, it attracts government attention.

Steve: Well, and we have Part 2 of Aaron's question.

Leo: Coming up next. I was thinking about going to - and the other reason is many people have bandwidth restrictions. There's, you know, a home VPN server isn't always a perfect solution.

Steve: Good point.

Leo: Yeah. Part 2 from @vaderIRL. [Darth Vader voice] I was just thinking about going to the Pi route. My father - we know who his father is. My father travels to Kuwait - oh, sure, Kuwait - for work and just needs a basic U.S.-based connection

from time to time. Well, there you go. That's a good way to do it.

Steve: Yeah. So anyway, so this was in a little Twitter dialogue that I had with him. And so I sent back, I noted, I said that's perfect. The only glitch is that home IPs can drift. So arranging some DynDNS is useful for finding a router's current public IP. And DynDNS is short for dynamic DNS. And the idea is that it's a service, a public service, a publicly accessible service that your router informs of its current IP so that you're able to query the public service. And then it will tell you, it will inform you of any changes to your router's IP.

I did a little bit of looking around, and I remembered that there were some registrars that offered that as part of their package. Unfortunately, Hover doesn't. Hover is my chosen registrar. Whenever I can get a domain from them, I do. There are some top-level domains they still don't support, so I'm stuck on a couple of those with Network Solutions. But I'm really happy with Hover. But they do not offer a dynamic DNS natively. Namecheap does, and Google, the Google Domains service does.

And it looks like Google Domains, for a dotcom domain, is $12 a year and includes dynamic DNS support. So many people already have a relationship with Google. So to me that seems like the path of least resistance. If you were interested in doing this, make up some crazy domain for yourself. You will have a dotcom domain, maybe for the first time in your life, your own. Host it with Google, and for a dollar a month Google will provide the domain registration and supports dynamic DNS so that you're able to find your router using your own custom dotcom domain wherever you are out in public. And so this is...

Leo: There are routers that do this, too. My Asus router does DynDNS. A lot of routers will do that.

Steve: Correct. Although you need to use a third-party service.

Leo: Yeah, they offer it through their [crosstalk].

Steve: Yeah. And there were - it used to be free. Now they're beginning to charge because they're...

Leo: Ah, that's too bad.

Steve: And so it turns out, it looks to me like Google is the - I wasn't able to see anything that was as good as 12 bucks a year, a dollar a month for that service.

Leo: Right. Yeah, that's a good deal. What was I going to say besides that? Oh, also your router may have a simple OpenVPN solution. You can kill two birds with one stone.

Steve: Yeah.

Leo: Darth Vader. Scott Pritchett asks via Twitter, he's @bitman: I don't understand. I don't understand, Steve, how a memory page could be identical between VM instances, yet contain their private key.

Steve: So I don't think I was as clear - he's referring, of course, to last week's coverage of the Flip Feng Shui exploit. And I wasn't as clear as I should have been. I did say that it was the public key. But I'll say it again, it's the public key, which is available because it's freely given out as part of the authentication. And this is something I don't think I've ever really explicitly explained. I've shied away from it because you know I don't like to be inexact. And the math is very tricky. But here's the way we can think of it. The public key contains the private key because the private key - and I'm deliberately simplifying. But the private key is one of the two primes.

So essentially what we know is that it's impossible to separate the multiplied primes in reasonable time. So the public key contains the private key. And the trick is that you can't pull it back apart. So essentially it's hiding in plain sight. The public key has the private key as part of it. And it's only if you knew the other prime that you could then divide that by the public key in order to extract the paired prime which is the private key.

So now you can understand why this bit-flipping works. If you do something to it that suddenly makes it much easier to factor, then this problem of deliberately created non-factorizable huge primes disappears. But that's really the key to the way this crypto works is that the private key is one of the two primes that is multiplied to give you the public key, but no one knows what it is. There's the public key. Somewhere in there is the private key, but you can't find it. It's literally hiding in plain sight. Which is so cool.

Leo: Here you go. Steve Gibson. You ready for more?

Steve: You betcha.

Leo: The action continues with Question 5. Stanislav Leaderman in Oregon - I love this question. He wants to know how to be secure with less technology. With less technology. I used to be a regular listener and strong follower of Steve's suggestions related to security. I bought SpinRite seven or eight years ago. I still like it to this day. However, over a few years I haven't been focusing so much upon security due to other distractions. Moreover, security has become so much more complex that an ordinary person can't get his head wrapped around what one should and shouldn't do.

For example, banks require your phone number to send a text message. Yeah, that's a good security measure. But then they sell your phone number to telemarketers, and you get sales calls late in the evening. I used to give banks my Google Voice number, but then they began requiring a physical number. I just went through an ordeal with CITI credit cards. They wouldn't accept a forwarding number like Google Voice, and they knew Google wasn't mine since they rely on some third party who apparently has all my personal information, including my actual cell phone number.

Steve, what are the prudent measures one could take to secure accounts like banks and emails against fraud, but also not to be so caught up in this that it doesn't take

over and become your entire life? Sorry for the lengthy email. What measures would you recommend to implement? Paper passwords? YubiKey? I have LastPass, but not a smartphone, only a flip phone. It seems many applications like LastPass need a smartphone for their multifactor authentication. Thank you.

Steve: I thought that was a really great question because we're so steeped in technology and our smartphones that it's easy to fail to appreciate that some of these solutions aren't applicable to everyone. And it's like, yeah, it's easy to say, oh, every website you use needs to have a unique 20-character crazy high-entropy password. But, boy, I mean, I don't know how I would do it if I didn't have LastPass or a similar password manager to manage that for me. And we've been talking about how SMS second factor is falling by the wayside and now formally being deprecated in favor of the time-based password. But that requires an app on your phone, which is, you know, we just assume someone's going to have a smartphone. But what when you don't?

Leo: There are desktop authentication apps, though. You don't have to have a phone, only if you're mobile.

Steve: Right.

Leo: And you can always have it texted to you. Well, not always, but in many cases.

Steve: Right. So anyway, I did appreciate his observation. I mean, this is really becoming a mess.

Leo: Yeah, yeah.

Steve: And smart people are trying to come up with solutions, and we keep trying things. I don't have a great answer. And especially for somebody mobile who is unable to use the technology which sort of does, I mean, I would argue he's right. It requires a smartphone.

Leo: Right.

Steve: If you've got a flip phone, well, you can't play.

Leo: That's the nice thing about authenticator apps. You don't need to give them the phone number. Although you're going to give the bank your phone number anyway for all sorts of other reasons.

Steve: Good point.

Leo: Get a bank that has a good privacy policy. If they, I mean, I wouldn't assume they're selling your phone number. We all get calls all the time. I have an unlisted number I never use, but I get calls on it, soliciting calls, because they're calling random numbers.

Steve: Yeah. I have two physical landlines, and I get the same robot spaced about 30 minutes apart on each phone.

Leo: Yeah, it's an automated dialer.

Steve: Because they're just - yup. That's all they're doing.

Leo: At this point, they're not targeting people anymore. They just call everybody. It's so cheap to get labor in India and elsewhere. You know, Lisa and I were driving to the football game last night, and a call comes in from a very long international number. And I thought, oh, maybe it's Abby. I'd better answer it because she's in Mexico. And sure enough, "Hello, this is Microsoft. We've been seeing some unusual behavior on your Windows system. Are you in front of your computer right now?"

Steve: Yeah, please launch event viewer. And, oh, look at that, yeah.

Leo: And I thought, I could play along with him. And instead I said, you know, "Shame on you for scamming people. You should not be doing this."

Steve: Good.

Leo: "How does your mother feel about you doing this?" I tried to humiliate him. I don't know if he - he hung up, by the way. I don't know if he...

Steve: Jobs are scarce, I think.

Leo: Yeah, and I understand. But then scamming people is not the solution. And by the way, that was a cell phone number. It wasn't my Google Voice number. It wasn't a public number. It's not widely available. They're calling random numbers.

Steve: Yeah.

Leo: You know. So don't assume that your bank has sold your number. Bu if they do, if you really think they are, get a better bank.

Steve: And caller ID is now spoofed. I noticed...

Leo: And that's why you can't tell.

Steve: ...that things show as my area code.

Leo: Oh, yeah.

Steve: It's like, okay, now.

Leo: Yeah, yeah. It's always my area code now. "Hello. Your auto insurance is about to expire." No, it's not. "We have a very important offer for Leo Laporte." Actually, they don't usually know my name. It's, you know, I wish there were some better solution. But certainly, look, ask the bank what its privacy policy is. And I would never, I would never patronize a bank that is selling your information of any kind to anybody. Banks have to have a lot of personal information about you. You want to be able to use your phone number with a bank; right?

Steve: Yeah. They are fudging in that direction, though. They send you their privacy updates in fine print. And it's like, oh, okay, just do whatever you're going to do.

Leo: Yeah. Michael Zimmermann, Sydney, Australia wonders how SQRL can evolve - into a killer SQRL. No, not that kind. Hi, Steve and Leo. There are classic engineering designs. We all remember the 1963 Jaguar XK-E Roadster. My best friend in sixth grade's dad had one of those.

Steve: Ooh, yes.

Leo: Yeah. The 1955 Mercedes-Benz 300 SL. That's the one with the doors that went - the gullwing doors.

Steve: Like you have now, Leo.

Leo: I have falcon wing, please.

Steve: Ah, falcon, upgraded from a gull.

Leo: Let's suppose you have one, but you can only keep it if you convert it by swapping out the engine with the latest electric model. Quite a challenge.

Now, I started thinking about SQRL, how it is going to be a classic design with very high adoption across major websites. That's what we all hope for. As I understand it, you have your master key and your web URL as inputs, the black box engine where

the crypto runs, and the resultant output, which is used by the website to identify you. You have explained how you can rekey your master key.

But I was wondering what happens if your engine is no longer safe, and you need to swap it out for a better model. A different engine would produce different output; right? Does the client and server have the ability to say we are using crypto engine V6, V8, electric, warp drive? Or Mr. Fusion? Thanks for the show, and looking forward to your explanation. Please let us know how to push SQRL - I think he means push it as in get the word out - once it's released.

Steve: I hope that's what he means.

Leo: I think that's what he means.

Steve: Yeah. So the answer is yes. And it's as simple as a version number, which is built into the specification. And the way it's - remember what we've done with SSL and TLS, "we" meaning the industry in this case, where it was always built with a large array of different suites that it could use, where the browser would offer the ones it knew about, and the server would respond with what it understood. SQRL has a similar facility. It's a compound version number which can contain any information. Right now it's set to one.

But, for example, if it turns out that there is something where we need to evolve, if any fundamental change needs to be made, either the crypto needs to be evolved, or a problem is found, or we want to add some feature, not change one, but add it to a future one, then the client could say v1,2 or anything it wants to, really, depending upon how the spec evolves. The server then receives that, knows what the client understands, and then the server's response is which version it chose of what the client knew. So very similar to that model, and that protects us and makes us future proof. So, yup, we got that.

Leo: Question 7 comes from an anonymous listener who shared a brief note about Windows Update issues. He says: Can you address this for us, Steve?

Steve: Yeah. This was just - I put this in here as a reminder to me. Thank you, Leo.

Leo: This is Woody on Windows. You've quoted him before, Woody Leonhard.

Steve: Yeah. And I guess Windows 7 is having a real problem with getting updated after a clean install. Some people have said that they imagine the reason that Never10 is still being downloaded, because I've built-in that little grabber of Windows Update Update, and it just makes it easy to apply. What's very cool, though, I did some digging, and I found an answer which is very comprehensive at answers.microsoft.com. The link's in the show notes, but it's also made it now into the permanent link database which is the link farm, GRC's link farm. So anyone - and believe me, this is what you want to use. Even better than the Windows Update Update Update Update link that I have already in the link farm, it walks a Windows clean new install through a surprisingly daunting process that involves interrupting normal things, setting it to never update, rebooting it, going in

and stopping a service.

Anyway, this guy has worked through what it takes to absolutely nail it. So I just wanted to let our listeners know. I know that we have many people who are having whatever occasion to install Windows 7, maybe only for testing purposes in a VM, if not in a new system. But from everything I've seen, this really works. It's new. It's mid-August, I think, was the date of it. And if anything changes after today, actually, Patch Tuesday is today, then it'll be updated. But if anyone's curious, GRC.com/linkfarm in the upper section of permanent links, it now has a home there because it looks like this is the way to do it.

And in fact the guy's posting starts out saying: "Windows Update has become quite problematic for Windows 7 users. For the past year or so, we've been working to find a solution that will work for you. We have found one that works very well indeed for most. We know for certain this works well for August 2016 until this September 2nd Tuesday, known as Patch Tuesday. We hope to be able to update this for September." Anyway, anybody, I would say, in the future check this posting because they will keep it updated, and I wanted to make sure people could always find it.

Leo: Nice. GRC.com/linkfarm.

Steve: And, boy, you know, I guess it's unfortunate that it's that difficult. I can't imagine Microsoft did this deliberately. I think it's just that the original Windows 7 image, which is SP1 is the latest image you can start with, it itself is so old that what Microsoft didn't do was provide backwards compatibility for Update all the way back to there. And so you sort of have to bootstrap yourself. I mean, they provide the means to get there, but it's just not automatic. It's not install Windows 7, and then click "Give me all the updates."

Leo: Question 8 is a follow-up to last week's Flip Feng Shui episode. George Mallard in Texas wonders whether DRAM ECC could prevent Rowhammer. He says: I heard your podcast on Rowhammer. I have a quick question. Would parity error checking, or ECC RAM on a server, pick up that bit that was flipped and potentially flip it back? If so, do high-end cloud servers employee ECC RAM? Thanks for your show and SpinRite. I use it monthly on my system.

Steve: So that's a very good point, and it's something I skipped over in the coverage of Rowhammer last week, although the guys that wrote the paper did cover it extensively. In their Section 6.1.1 of the paper they wrote - and that section is about hardware, that is, hardware remediation. They wrote: "We recommend DRAM consumers perform extensive Rowhammer testing to identify vulnerable DRAM modules." And I'll talk in a minute about how end-users can do that because it turns out we can.

"These DRAM modules should be replaced; but, if this is not possible, reducing DRAM refresh intervals, for example in half, which I did talk about last week, may be sufficient to protect against Rowhammer. However, this also reduces DRAM performance" - but not by much, 1 or 2%, they wrote - "and consumes additional power. Another option is to rely on memory with error correcting codes (ECC)" - which actually works in a very similar fashion to the way we've often discussed it works on hard drives, where there's additional information stored with the memory that allows it to fix its own problems, essentially - "which protects against single bit flips. Unfortunately, we have observed," they write, "that Rowhammer can occasionally induce multiple flips in a single 64-bit

word confirming the findings of the original Rowhammer paper."

So not only did the original Rowhammer paper mention that, they've seen it in their own subsequent research. "These multi-flips," they write, "can cause corruption even in the presence of error correction memory. More expensive multi-ECC DIMMs can protect against multiple bit flips, but it is still unclear whether they can completely mitigate Rowhammer.

"A more promising technology," they write, "is directed row refresh" - which I did talk about last week - "which is implemented in low-power DDR4" - which is abbreviated LPDDR4 - "and some DDR4 implementations. Low-power DDR4 counts the number of activations of each row; and, when this number grows beyond a particular threshold, refreshes the adjunct rows, preventing cell charges from falling below the error margin." So it's cool, it's sort of a demand refresh-based RAM hardware.

"Newer Intel processors support a similar feature for DDR3, but require compliant DIMMs. While these fixes mitigate Rowhammer, replacing most of current DDR3 deployments with low-power DDR4 or secure DDR4 DIMMs is not economically feasible as it requires compatible mainboards and processors." So you're not just swapping the memory out. You need to change everything, essentially. "As a result," they write, "a software solution is necessary for mitigating Rowhammer in current deployments."

Now, what's cool is that the free, well-known, venerable Memtest86 has, since v6.2, added a Rowhammer vulnerability test. And I got a kick out of noticing, as I was doing some research, some people asking the question, "Hey, why does my memory pass all the tests except Rowhammer?" It's like, gee. I wonder why?

Leo: Oh, my.

Steve: Anyway, so Memtest86.com, M-E-M-T-E-S-T-8-6 dotcom. I imagine that'll be of huge interest to our listeners, who are just curious whether they've got Rowhammer-susceptible DRAM. You can now find out just by running that on your machine.

Leo: I'm amazed that program is still around. And not only still around, but being updated.

Steve: Yes, yes. It is still current. And I used it, oh, I used it - I use it in a perfunctory fashion whenever I'm building a new system, just as part of what I do, is after the hardware is there - because you boot it. It takes over the system, much as SpinRite does because it needs full access to everything. And it just - it's got a nice little text display, and it just cranks away and exercises your RAM.

Leo: It's amazing.

Steve: And you remember that I was having - it was a challenge for me. I was trying to use 128GB at high speed, and I was never able to get that to work. So I either had to drop to 64GB at high speed or go to a lower speed to get 128GB. And I opted for high speed and 64GB. Although I've not done any extensive benchmarks yet. But I was using Memtest86 to essentially torture-test the RAM at all of these different speed settings as I

was setting things up. So it's a super useful utility.

Leo: Wow. Lai Min-Hui in Malaysia wonders about lithium-ion battery advice: For years, I've followed your advice that lithium-ion batteries loved to be fully charged, so plugging in the charging cable whenever we can is the best thing to do in terms of prolonging battery life. However, assuming that lithium polymer is equivalent to lithium-ion, Father Robert mentioned in TWiT TNSS-69 - that's The New Screen Savers #69 - that these batteries prefer to stay at 50% charged. I don't know who to believe anymore. Hope you can clarify. Well, they're both right, Lai Min-Hui.

Steve: Exactly. The good news is we both get to be right on this one. Here's what's important to understand. Not loving to be deeply discharged is not the same as loving to be charged. So what I have often said is that you do not want to run lithium-ion batteries all the way to the ground. And I've seen real-world evidence of this. I've got a good friend who's killed a number of his non-battery replaceable iDevices because his habit was just wait till everything turned red, and there was like a little sliver, a little itty-bitty sliver, and it said 2% left. Then he would plug it in and charge it up. And what do you know, six months later it was dead. It would no longer hold a charge. And I explained to him, I said, "No, no, no, no, no. Lithium-ion cells really do not like to be run to the ground."

So my advice about plugging them in is to keep them off of the ground. The challenge is that they also don't like to get overcharged. And so when you're right up there with the battery topped off, you're skating on thin ice because, if the charging technology is not good, it can push it too far, and then you start damaging the battery. So, and then we were just talking about this a couple weeks ago, actually, how impressed I was that Lenovo noted that I had my Carbon X1 third-gen laptop plugged in. That's like, it was just living on the adapter. And it popped up a little balloon and said, hey, we see that you seem to be, like, just docked all the time. In that case, bringing the battery down to half-mast will make it happier. And so I said, oh, that's brilliant for you to, like, realize that. And so I gave it permission, and it let the battery come down, and now it holds it at 50.

So the problem, of course, is if I then grabbed it to take it out on the road, it's already half discharged. So the idea is that, in this mode, I would turn that off. It would top off the battery. Then I would unplug it and take it on the road with me and use regular cycling. But when it's living on the adapter, much safer and better for the battery over the long term if it's kept at about half full. So that's what Father Robert was saying.

And I did want to take this opportunity to draw the distinction between "love being charged," which they don't quite so much. What I really meant, though, was them really disliking being deeply discharged. So the best behavior is - and Apple seems to have nailed charging technology. I don't have any problem with any of my devices living on their charger. Every one of them is plugged in all the time except when it's not, which is typically just briefly. Like I'll take my phone out with me, and then when I'm home it gets plugged right back in.

Leo: So keep it plugged in. It would be nice if they had a setting that said "keep it at 50%."

Steve: Right. And in fact, Leo, you should know, before my Palm Tungstens went into the refrigerator, I did bring the charge down about 75%, and then I disconnected the

battery inside so there would be no long-term leakage, and then they went into the deep freeze.

Leo: Oh. Will they come out at 75? There's no leakage at all?

Steve: No, they will probably be discharged. But it's better if they drop from there than if they drop from 100.

Leo: Right. It's hard for me to imagine a scenario where you would actually take your Tungstens out of the freezer for use.

Steve: And you know, Leo, when I look at that sad puppy now, I think, why did I...

Leo: What was I thinking?

Steve: Why did I ever think I would - well, you know, I was not wrong about my HP calculator. I've got nine of those. I keep waiting for one of them to die. But I just - I never want to be without this calculator. Although now, 42, you know, PCalc.

Leo: Right, excellent.

Steve: On iOS? Oh, it's - I use it whenever I'm not in front of my physical calculator. So some things don't get better; some things do.

Leo: How much time do you actually use a calculator, though? I mean, I haven't used a calculator in years?

Steve: Oh, I'm - constantly.

Leo: What are you using it for? Balancing your checkbook?

Steve: No. Engineering.

Leo: Oh.

Steve: I do a little electronics on the side.

Leo: Engineering, yeah. That's what they call it now.

Steve: I'm looking for a sample, but my engineering pad's over in the...

Leo: Are you - oh, okay. Like you mean the Portable Dog Killer kind of engineering.

Steve: That kind of stuff, yeah. Circuit design.

Leo: Circuit design. Yeah, well, that makes sense. Sitbit in London, Ontario, Canada has our last question of the day. He weighs in on the debate stimulated by television's "Silicon Valley" on tabs versus spaces. You showed us last week a graph. Somebody had analyzed hundreds of thousands of GitHub submissions and came up with a winner. Sitbit says I think what the Googler actually discovered was the default setting of the preferred editors used for the various languages because, well, for instance, anyone writing C is probably also using VI.

Steve: So I thought that was a useful and interesting observation. I would argue a little bit - okay. So first of all, I would absolutely give him credit, him or her, Sitbit, credit for observing that there would certainly be a bias. Although I can't think of any population more than programmers who would tend to change the settings to what they want.

Leo: Right.

Steve: So I've often spoken of the tyranny of the default, how browsers typically are just going to be left the way they're set by most users. I think programmers and programming editors probably customize to a much greater degree. Although I did want to - I wanted to use this as a reminder and trigger that several people wrote to let me know that the reason why Golang has almost 100% tabs, as was shown on that chart - remember I noted that it was Go and C that were the two strongest tabbed languages, is because Go contains a built-in formatter, go-fmt, which is always run. And it enforces, it programmatically enforces the use of tabs. So tabs, there's no question of tabs versus spaces with Go. You're using tabs because the system says, what are all these spaces? We can just turn that into one tab and...

Leo: Save space.

Steve: ...save seven bytes.

Leo: Yeah. Well, there you go. We weigh in on every possible important topic in the [crosstalk].

Steve: The definitive statement.

Leo: Steve Gibson's at GRC.com. That's where you'll find SpinRite, the world's best hard drive maintenance and recovery utility. That's his bread and butter. Everything

else, though, is free, including SQRL. You can find out more about his login solution, his Perfect Paper Passwords. He mentioned those. ShieldsUP, very famous for that.

Steve: And even the passwords page, an amazing number of people just go there to get random gibberish every day.

Leo: Well, of course. What a great tool.

Steve: Yeah. I've got good gibberish, Leo.

Leo: Good gibberish. It's truly random, not pseudorandom. None of that pseudorandom gibberish. No, it's real gibberish, all right. And the sleep formula, which I plan on using on Sunday.

Steve: Great.

Leo: When I will be massively jetlagged after flying to Europe. Although somebody has sent me, I haven't received it yet, a book, apparently long out of print, but highly esteemed by global travelers - you could buy it on eBay for 100 bucks - with the ultimate jetlag cure. It involves carefully dosed bits of caffeine and peanuts and sleep and fake sleep. I haven't seen it yet.

Steve: Wow.

Leo: But it sounds like a highly complicated but surefire cure.

Steve: Oh, that'll put you to sleep.

Leo: Just read this book. So I wish there were some way I could do a kind of clinically controlled test comparing Steve's ultimate sleep formula - because that would be good for jetlag; right? I would think.

Steve: Which?

Leo: Yours. Because what I'll do is I'll stay awake best I can till the sun...

Steve: Then knock yourself out.

Leo: ...goes down and it's nighttime. The problem, of course, is it'll be nighttime in

Paris at 8:00 a.m. here. And so my eyes are going to go boing, it's time to get up. That's when Steve's Healthy Sleep Formula will kick in.

Steve: Yeah. I updated the page yesterday with the news of the third ingredient.

Leo: Oh. Oh, dear.

Steve: And I didn't anything because I wanted the people - 2,700 people a day go to that page.

Leo: Geez, Louise.

Steve: I don't what is happening. But the problem is, and I mentioned it quietly on this podcast about a month ago, it's something called oleamide. And some researchers discovered it in the cerebral spinal fluid of sleep-deprived cats. We have it in ourselves, but of course no one wants to have their cerebral spinal fluid tapped. So the cats got used for that. Anyway, so it's a natural compound. It nails the formula. And I'm only talking about it now because it lasted about 12 hours, and now they're all sold out of that.

Leo: Oh, man.

Steve: I know. So I had written them. On Sunday I wrote them a long note - the company's called LiftMode are the people in Chicago who offer it - and telling them what was going to happen, giving them a heads-up, asking them how much of this can you make? Because you're going to have to start making a lot of it because it perfects the formula. And I've just - I'm been biding my time because I wanted to get this one right. And I did also talk to Source Naturals. They promised me that niacinamide will be back at the end of the month. So that's finally going to be available again, after the entire planet sold out when it became part of the v2. So anyway, I have not heard back from the LiftMode people about oleamide. I've got my fingers crossed that we won't have to wait long.

Leo: Looks like they're out of stock, though.

Steve: Yeah, I know. I refreshed the page before the podcast. They had 10 available a little bit, like an hour before the podcast, and now it's gone.

Leo: Oh, well.

Steve: Yeah.

Leo: Maybe I'll have to do peanuts and coffee after all.

Steve: Well, no. You have the two components. You have the niacinamide and the melatonin.

Leo: Yeah, yeah.

Steve: And that does a good job. It wasn't enough for me. The oleamide - oh, and the other thing is that the two-piece formula didn't solve sleep initiation. It only was targeted at sleep maintenance. Well, the oleamide knocks you out. So it also solves the sleep initiation problem.

Leo: So our flight leaves at 8:00. I was thinking I'll have dinner, watch a movie. Around 11:00 I'll go to sleep.

Steve: Didn't you say you were going to go and see the IPTV guys? I mean the...

Leo: That's next. That's after we get back. They're in Gainesville.

Steve: Oh, okay.

Leo: That'll be another one. Lisa calls them "turn and burns."

Steve: That's right.

Leo: Yeah. So, okay. Well, you know, I might just be tired. That sometimes happens. I can live with that.

Steve: Ultimately, I think the way to solve the jetlag problem is just go to sleep.

Leo: Yeah.

Steve: Sleep as long as you possibly can and just...

Leo: Yeah, yeah. It's a good - it's nature's natural nurse.

Steve: Yeah.

Leo: Ladies and gentlemen, if you don't go to GRC.com to get the sleep formula or other stuff, or SpinRite, you could go there to get this very show. Steve has audio versions and transcripts of the show. We also have audio and video. No transcripts at our website, TWiT.tv/sn. And by the way, I don't care where you get it. It's all good. You can also get it in your favorite podcatcher because a lot of people subscribe to this show and then save it on their massive Drobo drive, all 577 episodes with Steve, because you never know when you're going to need some Security Now!.

Thanks, Steve. I won't be here next week. I think Father - you're doing a show Friday, right, with Father Robert, or something?

Steve: Well, we're not going to see you for four weeks.

Leo: No, no, no, no, no. October, you'll see me October - you're right. October 11th I'll see you.

Steve: Yeah. I believe Father Robert will be my co-host because we did arrange - he had a scheduling conflict. And so we're recording one of them at an odd time in order to keep continuity.

Leo: Yeah, I'm going to miss three episodes, yeah.

Steve: Our listeners will miss nothing.

Leo: No, because I'm not the point of Security Now!, Steve is. I am merely your amanuensis, your maitre d', your...

Steve: Facilitator.

Leo: Facilitator. I'm here just to make sure the pages are turned. But Steve will be back next week. We do it every Tuesday, 1:30 Pacific, 4:30 Eastern, 20:30 UTC, if you want to watch live and be in the chatroom. Otherwise, of course, download it anytime. Thanks, Steve. We'll see you next time. I'll see you in a month.

Steve: Thanks, Leo. A month it is. Bye.