Listener Feedback #230

Description: Leo and I discuss the week's major security events - including lots of new fur flying over the escalating Apple v. FBI/DoJ encryption battle - and discuss questions and comments from listeners of previous episodes. We 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-551.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-551-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We've got questions and answers and a little bit of security news, too. And we'll talk a little bit more about encryption, some interesting stories coming across the news, the security news desk. Security Now! is up next.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 551, recorded Tuesday, March 15, 2016: Your questions, Steve's answers, #230.

It's time for Security Now!, the show where we cover the latest security and privacy news with the Explainer in Chief, Steven "Tiberius" Gibson. I crack myself up. Hi, Steve.

Steve Gibson: I'm so glad you're having fun over there, Leo.

Leo: Well, it's because I had such a good night's sleep, thanks to you. I must confess.

Steve: Ah, yes.

Leo: I went back to the elephant dose, by the way.

Steve: Good. Well, that's, you know, take what works.

Leo: Yeah. Oh, it's nice, I mean, I'm just out. And then I looked at my - I wasn't using the Zeo forehead sleep tracker. I had a Fitbit on, and it's the first time I've - I should show you the graph. It's the first time I've ever had sleep where there was no - it was just, like, solid sleep. There's no - it was no waking moments, not even light sleep. It was just like a couple of tosses and turns at the beginning and, boom.

Steve: Yeah, I've heard from some people who remarked...

Leo: I thought it was broken.

Steve: That remarked that they knew they slept well because when they woke up they were in exactly the same position as when they went to sleep. They hadn't gone to the other side of the bed or migrated. The covers were not ruffled.

Leo: Lisa says, "You don't wake me up in the middle of the night anymore."

Steve: Anyway, I put a little note down in miscellaneous stuff for us to talk, to bring this up.

Leo: Steve Gibson's, what do you call it, Miracle Sleep Formula.

Steve: Well, and your comment, I watched you on The New Screen Savers talking about it and the FDA and my use of the term the Healthy Sleep Formula.

Leo: Yeah.

Steve: So anyway, we'll talk about that. We've got a lot to cover. This is a Q&A episode #230. That is, our 230th Q&A, but Security Now! Episode 551. And we're going to talk about, you know, I keep trying to think I'm - I guess I'm deluding myself to think that we are ever going to be able to stop talking about this encryption debate, dispute, whatever it is, because it's just not going away. And a lot has happened in the last week. We need to...

Leo: Well, and we expect you to talk about it, of all people.

Steve: Right, right.

Leo: The other thing, and someday I'd love to know, if you think it's possible to make a smartphone that is truly secure, regardless of government interference, like Blackphone or something, if it's possible, even.

Steve: Well, your point, the point you raised a couple weeks ago was, I think, very salient, which is one of the problems we have is that this is a very complex topic. And it's easy to get tunnel vision on one specific aspect, like, oh, how long is the key? And how strong is a PBKDF algorithm so that it's 80 milliseconds per guess, blah blah blah? And then someone says, yeah, but if Apple installs a keystroke logger...

Leo: Boom.

Steve: And it's like, yeah, all bets are off.

Leo: Boom.

Steve: So the point is...

Leo: Even if you're using WhatsApp or Threema; right?

Steve: Correct. Correct. If you put something upstream of all of this unbreakable, industrial-strength, military grade, even the NSA, you know, they disapprove of it because they can't get into it, it doesn't matter. If you capture their keystrokes or you record their microphone before it's encrypted, then it's all bypassed. So I think it's necessary. So the point you always raise when, for example, I talk about a given system that I'm enamored of, like Threema, for example, is, well, yes, but we have to trust them.

What I think has come out of this, and I haven't seen this written anywhere, or noted, is Apple is demonstrating the degree to which they're serious about defending our privacy. And if the FBI thought they were getting some PR benefit from encrypting their phones, then the FBI has just given them a massive shot of adrenalin and help by attacking them in this way because you can't turn around. I mean, the servers in restaurants are talking about it.

So anyway, so, yeah, we have - there's some more interesting news because the government has responded in a footnote that just - I'm beginning to see red when I hear some of these things. But then there's also some news that's happened much sooner, it was reported on The Hill publication yesterday, about Dianne Feinstein and a senator from South Carolina. Anyway, we'll get to that. Of course John Oliver has - that thing went viral from Sunday night.

Leo: I haven't seen it yet. He does his trademark talking about a subject, which he's done so well in the past. This time he's talking about encryption.

Steve: Yes. And you know that I'm not, as I've said to you, I'm not a big fan of his. Jon Stewart and Colbert, before Colbert switched over to CBS, I loved them both there. John, I just never got it. So I don't TiVo him. But I got so many tweets that I thought, okay, finally. So yesterday afternoon I watched it all the way through. It's 18 minutes long. It is a masterpiece. And so I tweeted out to my followers in case they hadn't picked up on it, you've got to see this.

Also I wanted to talk a little bit about David Pogue's comments on TWiT because he asks really good questions. And I blogged - for the first time, someone commented, since 2014 - on this issue. We have a new IoT nightmare, I mean, a specific IoT threat. Some news about BleepingComputer getting sued. A new horrifying DDoS attack multiplying vector. Some news that I know you've been covering about Windows 10 pushing it even further and harder.

Leo: You're starting to win me over with your point of view on Windows 10. Well, I should say Microsoft is starting to win me over to your point of view.

Steve: Yeah, yeah. And then I do have a Windows Update Update Follow-up Follow-up, and some miscellaneous stuff, and then we're going to do 10 questions from our listeners.

Leo: Wow.

Steve: And they're great ones. So I think we've got a good, engaging podcast. And everyone, just grip the steering wheel firmly as you're driving, if you're listening to this while you're commuting, because you want to stay in the center of your lane, if possible.

Leo: Stay in the lane. Now, somebody the car will do that for you. But until then you'll have to keep your hands on the wheel and your eyes on the road.

Steve: Right.

Leo: But that's why we do these in audio as well as video. You can listen, continue to listen. Speaking of guys that do a great job explaining this stuff, there's nobody better than the Security Explainer in Chief, Steven "Tiberius" Gibson.

Steve: So I found, for the Picture of the Week on the first page of the show notes, an interesting chart which Mozilla posted. On their Twitter feed they linked to it. They described it as "bending the curve for HTTPS adoption," which has quadrupled since Let's Encrypt launched.

Leo: It's kind of amazing; isn't it?

Steve: It really is. Now, looking at it, there are a few things to note here. First of all, it's the rate of adoption which has quadrupled. They didn't make that clear. So it's not - it didn't jump. The other thing that they did, and, I mean, I understand they had to, but it's one of my pet peeves, is when you use...

Leo: I know, no zeroes.

Steve: Yes, non-zero-based graphs.

Leo: Yup.

Steve: So this looks like - it looks awesome. And then you look at it, and wait a minute, on the Y axis it goes from 37 to 42.5. So it is vastly less dramatic. On the other hand, it doesn't change the fact that the ratio is 4:1. So in this case it's okay because they're not lying through graphs. In fact, if they did use a zero-based graph, both lines would look almost horizontal, and it wouldn't be as clear that the second half at the Let's Encrypt launch point had quadrupled its rate of increase. So anyway, still, what they were seeing was, before Let's Encrypt, was, yes, sites were generally moseying along, slowly adding encryption over time; and that rate is four times what it was before.

So a really nice data point, in addition to the fact that we know more than one million sites in - what was it? We reported it. Since it was near December, it must have been the first three months, they got a million sites, or a million certificates and 2.5 million domain names covered because of course certs can have multiple domains in them.

Leo: One of our chatters, though - this is bad news. David Johnson in our chat says Apple's iTunes servers do not accept Let's Encrypt certificates. So his podcast isn't updating in the iTunes store.

Steve: That's odd because they cross-signed them with an old-school major CA. I don't remember now who it was.

Leo: He says they're using Java 6 for some of their components. So it must be a local client.

Steve: Yeah, I think it's their fault. I wouldn't blame Let's Encrypt.

Leo: Oh, I'm sure it is. Yeah, yeah, yeah, yeah, yeah.

Steve: So what must be is that all of us have a root CA which includes the one that is cross-signing the Let's Encrypt certs. But this particular little weird client lacks that root. And what I would do is I would update that client's root. You would certainly want to keep that up to date.

So, okay. There are a number of things about the whole encryption debate/debacle. First up, I wanted to note that the rhetoric, even the legal rhetoric keeps escalating. And I'm a little disappointed in what I'm seeing coming back from the government. I mean, they have not, they've certainly not been light-handed in this. There was a 45-page DoJ filing responding to Apple's response to theirs. And in a footnote they commented - let's see what I wrote here. I'll just quote myself.

"The latest filing in the legal war between the planet's most powerful government and its most valuable company" - this actually was Reuters who wrote this - "gave one indication of how the high-stakes confrontation could escalate even further. In what observers of

the case called a 'carefully calibrated threat'" - and, see, this is where I'm objecting to the idea that the DoJ should be threatening Apple - "a 'carefully calibrated threat,' the U.S. Justice Department last week suggested that it would be willing to demand that Apple turn over the source code that underlies its products, as well as the so-called 'signing key' that validates software as coming from Apple."

Leo: That's doing them a favor. We'll write it, if it's too hard for you.

Steve: Precisely. Apple is saying the All Writs Act is being overstepped because we know that software is free speech, and the government has no right to compel us to write code on its behalf that we don't want to write, that we explicitly believe we should not write. And so the DoJ is saying, okay, fine, we'll write it. Which means give us the source code, which we'll use to edit. And we'll of course need your super secret private signing key in order for the software, the version of iOS that we write, the DoJ OS, to be accepted by the iPhones. So, I mean, this is just - it's becoming kind of silly.

Leo: Although I remember China and India asking BlackBerry for the source code for the Exchange Server so that they could write a backdoor for BlackBerry Exchange Server. Right?

Steve: Right, yeah. And we covered that years ago, where BlackBerry was - I'm trying to remember the details. I think they wanted the servers to be relocated into their country, which would then give them access to the otherwise encrypted communications. And if the servers remained in Canada, where they were at the time, then they had no access because it was encrypted from point to point. But if it went through the BlackBerry servers, they believed that would give them some access. And I don't remember how that ended up.

Leo: Some of them got it, and some of them didn't. I think India got it, and BlackBerry was able to successfully fight it elsewhere. You know, I think what's really become clear, though, is the government's long-term goal is to have all encryption have a backdoor, period. And if that requires rewriting firmware, so be it. I mean, that's really what they want, isn't it.

Steve: It's why I abandoned CryptoLink, remember, a couple years ago.

Leo: They're going to outlaw, I fully expect them to outlaw encryption that doesn't have a governmental backdoor in it.

Steve: And here's what's shocking is how quick this is going to happen, which is the story after the next one. I wanted to mention first that all of this Apple issue has been the data-at-rest question. That is, law enforcement handing Apple a phone and saying, "We need to get into this, and here's a court order authorizing you to lawfully open this phone for us." Now, the other side of this is the data-in-flight question. And that's a whole different issue. But now WhatsApp is in the DoJ's crosshairs.

The New York Times summed it up beautifully in just the first couple paragraphs of their

story. They wrote: "While the Justice Department wages a public fight with Apple over access to a locked iPhone, government officials are privately debating how to resolve a prolonged standoff with another technology company, WhatsApp, over access to its popular instant messaging application. No decision has been made, but a court fight with WhatsApp, the world's largest mobile messaging service, would open a new front in the Obama administration's dispute with Silicon Valley over encryption, security, and privacy.

"WhatsApp, which is owned by Facebook, allows customers to send messages and make phone calls over the Internet. In the last year, the company has been adding encryption to those conversations, making it impossible for the Justice Department to read or eavesdrop, even with a judge's wiretap order. As recently as this past week, officials said, the Justice Department was discussing how to proceed in a continuing criminal investigation in which a federal judge had approved a wiretap, but investigators were stymied by WhatsApp's encryption."

So this is all, like, really pregnant. I mean, this is really coming to a head. And then I read just this morning on The Hill, they have sort of a "week ahead" summary. And the headline was "Senators close to unveiling the Burr-Feinstein encryption bill."

Leo: [Growling]

Steve: I know. Our California...

Leo: Our senator.

Steve: Senator Dianne Feinstein. And I watched her being interviewed, and it just - I just get tense because she is so far out there in the "no question whatsoever, nothing should be outside of the government reach" side. And it's just like, okay, calm down. Anyway, so she and North Carolina's Richard Burr are - they have legislation. It is written. And this week it is expected to be put forth and unveiled. And then of course this would be new law. Then I assume that both houses of Congress have to look at it and then reconcile it and come to an agreement. Then it goes to the President. And the question is, based on his SXSW comments this weekend, one wonders...

Leo: Well, we know he'll sign it. I think actually, you know what, I think it might be hard for Congress to pass a bill like this.

Steve: Oh, good.

Leo: I don't think it's going to get through Congress. I think Congress, actually, there's enough people on both sides of the aisle who are against this kind of thing that I think this might be - I'm hopeful that that - and by the way, that's the last stand because we know the President will sign it.

Steve: Yeah.

Leo: So sad.

Steve: Yeah. The Hill wrote: "While some argue that a judge should order WhatsApp to help investigators obtain the information they need in a readable format, others are hesitant to escalate the dispute, given that some lawmakers are expected to introduce" - oh, and this is about - "expected to introduce legislation to give law enforcement access to encrypted data as early as this week."

Now, the good news is, I mean, from a standpoint of, like, ripping the bandage off of the cut or the scab, I want to see what this legislation says, and I'm glad we don't have to wait four months or six months. I thought we were going to. But I guess this is like really ready. And I hope you're right, Leo. I hope that, see, I guess my feeling is we had all the demonstrations from the academic and the crypto community. We've had testimony that was sane and coherent. And unfortunately, if the legislation says what it must say, all of that's been ignored. I mean, the fact that this is coming from Dianne tells us that this is - the government can - well, I guess, okay, I'll finish that thought. The government can decrypt whatever they want.

Now, how do you implement that? I mean, we're all about details on this podcast. And as everyone knows, even if we do compel Apple to arrange to be able to comply with court orders, third-party apps, open source software, you know, existing crypto, which we've often said has already, it's like it's already left the dock. It's out of the barn. It's escaped. Everyone knows how to do that. So does it then become unlawful? Are you breaking the law if you use an encryption that isn't government-approved? Oh, this thing opens such a huge can of worms, I don't even know where to start.

So as we mentioned at the top of the show, I've not been a fan of John Oliver. I don't know, his style just doesn't hook me. But he did, and it was aired on HBO Sunday night, so two nights ago. It's on YouTube. You can find it easily. There's a link in the show notes. And I think in fact if you just put "John Oliver encryption," the first link I'm sure in Google will take you to their property in YouTube.

And even if you think you know everything about this, if you've heard everything we've said on this podcast, this is - it is a brilliant piece of work. And when I finished watching it, I found myself thinking, wow, they put a huge amount of effort into this. And there's also some fun at the end. So there is new stuff there, but really good coverage of this whole issue. And it's gone viral. Last time I looked it had a million-plus views, and that was yesterday afternoon after a bunch of people had tweeted me.

Leo: Three million now.

Steve: How many?

Leo: Three million.

Steve: Good, good. Although I don't know if it matters, except we know that legislators care what their constituents think. And these are three million constituents who have to come away with a better understanding, watching this, than they had before.

Leo: I think we're really seeing that Tim Cook's strategy was clear and was correct, which is to take it to the people, to create the conversation, and really as best we can educate people to the risks of a backdoor in encryption. And it is about backdoors ultimately now because that's clearly what the end game is for the government, and that's what the Burr-Feinstein bill is about and all that.

Steve: Problem is the term is ill-defined.

Leo: Yeah.

Steve: I don't like the term. So in my blog post at steve.grc.com, I wrote up a little essay on Saturday where the goal of the post was to separate the idea of encryption from access. And the whole point was one of the problems this discussion is having is that everyone talks about weakening encryption and backdoors and using terms that are not well defined.

The fact is we have the technology to allow Apple to give the government secure access on a phone-by-phone basis in a secure fashion. I've talked about it for the last few weeks, the notion of a completely random key which every iOS device has, which cannot be read from it, which is stored in the Secure Enclave, which Apple has encrypted under their ultra galactically secret private key so that, even if their database escapes, then it doesn't do bad guys any good because they still can't decrypt that database. The point is encryption does give us the tools to do this securely, if we wanted. And so the question is do we want that?

But my point was that needs to be a separate question. We weaken the argument for not doing it if we say we're not doing it because we can't, because no one believes us. And they're right, because we can. So it's important to separate the technology which can empower this from do we want to. And I think that makes it a much more pure and correct argument. And that's what David Pogue was asking was exactly this. Are we sure this is inseparable? And as I was listening to him on Sunday, I thought, well, I wrote about this the day before. That's exactly right. The technology will do anything we want it to. That part is true.

So it's wrong to say, and I think it weakens the argument to say, oh, we don't think we can. No, we can. The question is, is that what we want? And so I think that moves the argument to exactly the right place. Let's separate the technology, separate this encryption and all that. We could do it securely, if we wanted to. But Apple doesn't want to. That's the argument, I think, that we should have, is we don't want basically a spy state where the law is it's illegal to use strong encryption. I can't imagine being a citizen of that country, but it may be upon us. Anyway, so...

Leo: Again, and we'll probably have this conversation at some point, if encryption is made illegal - or real encryption is made illegal because anything less is not real encryption. But if real encryption is made illegal in the United States, will we be able to get hardware platforms that support it? And kids, go to ITProTV. Start learning security now because I said this 10 years ago, and it's starting to get more true than ever. The future freedom fighters won't carry guns. They'll be hackers. They'll be coders. They'll know how technology works. If you're going to fight for freedom,

that's how you're going to do it.

Steve: Like the first scene of "The Matrix," where the hackers are getting a little nondescript disk from Neo, and it's like this is all underground, super illegal stuff.

Leo: I can see a time when encryption, real encryption is illegal in the U.S.

Steve: Wow.

Leo: We've seen it before. It was classified as munitions in the '90s.

Steve: Yeah.

Leo: And we had a horrific consequence that we still suffer from.

Steve: Yup. I've been listening to you saying that, and you've exactly right. We've talked about it, how there's still 40-bit encryption lurking around because once upon a time it was mandated. And we've seen how slowly this stuff drains from the technical ecosystem. IoT devices, for example, many of them are never going to change after they come out of the factory. And they could have problems that will never get fixed. Well, they could have 40-bit encryption that'll never get updated because it was originally weak, deliberately weakened when it was released. And maybe that's what it'll be. I just - I can't wait to see this legislation. I'm so happy we won't have to wait very long for it because...

Leo: Oh, we haven't seen the text of Burr-Feinstein?

Steve: No. As far as I know it's not public yet. It's supposed to happen sometime this week.

Leo: Yeah.

Steve: But, boy, that'll be sad, Leo.

Leo: I don't - well, we'll see. I mean, I think the message...

Steve: Do we have to change the name of the podcast?

Leo: Insecurity Now!.

Steve: To Security Was!?

Leo: Security - oh, that's good, I like it, Security Was!. I like it.

Steve: Security Was!.

Leo: Yeah, that's good.

Steve: So speaking of Security Was!, Brian Krebs covered a story that came out of the Talos team. Cisco created a group called Talos Intelligence. And these guys, the Cisco team, took a look at a Trane - Trane is a very well-known, high-end HVAC - heating, ventilation, and air conditioning.

Leo: Oh, yeah, T-R-A-N-E, yeah.

Steve: T-R-A-N-E, exactly. But I think I had Trane equipment on my roof of the building that I built years ago, yeah, because it's high-end industrial. Although this is a consumer product because it shows you photos of your family on the color LCD. Anyway, so I'll just share this real quickly. Brian wrote, and I'm paraphrasing this: "Before purchasing an Internet of Things device - a thermostat, camera, or appliance made to be remotely accessed and/or controlled over the Internet - consider whether you can realistically care for and feed the security needs of yet another IoT thing.

"In April 2014" - so this is a little - this is dated, although it took Trane a long time to respond - "researchers at Cisco alerted HVAC vendor Trane about three separate critical vulnerabilities in their ComfortLink II line of Internet-connected thermostats. These thermostats feature large color LCD screens and a Busybox-based computer that connects directly to your wireless network, allowing the device to display, not just the temperature in your home, but also personal photo collections, the local weather forecast, and live weather radar maps, among other things.

"Cisco researchers found that the ComfortLink devices allow attackers to gain remote access and also use these devices as a jumping-off point to access the rest of a user's network." And this is like the nightmare scenario. This is the worst it could possibly be. "Trane has not yet responded to requests for comment. One big problem is that the ComfortLink thermostats come with" - everyone sitting down? Hold the steering wheel tightly and focus on the road. "ComfortLink thermostats come with credentials that have hardcoded passwords; and, by default, the accounts can be used to remotely log into the system over SSH."

Leo: Well, that's comforting.

Steve: Yeah, what a comforting link.

Leo: What a comfort.

Steve: Oh. So, and then "The two other bugs that Cisco reported to Trane would allow attackers to" - wait for it - "install their own malicious software on these vulnerable Trane devices..."

Leo: Oh, man.

Steve: "...and use those systems to maintain a persistent presence in the victim's local network."

Leo: Wow.

Steve: So the lesson here is there cannot be any question but that we need an isolated Internet of Things network. It's unfortunate, but just you cannot let this stuff - which is going to have all kinds of bells and whistles and be reaching out to remote web servers, thus opening up ports through your router in order to be contacted - you cannot let that stuff touch your main network. We have to have isolation.

And what I'm hoping, as I mentioned, is that the whole three-router solution, while it's simple and clean and robust, I'm hoping it's a stopgap, and that we will shortly have our standard consumer NAT router vendors explicitly offering network isolation for Internet of Things as a feature. That is, first of all, providing true isolation technology, not just saying it, but really doing it and making it easy to set up where, I mean, it should be easy. For the end-user it would be as easy as two different passwords, and you give your IoT things this password; you give your other network things that password. And by using different authentication, they are automatically on segmented networks that have no visibility of any kind to each other. Totally doable, just that easily. And we need that from router vendors, or maybe from third-party firmware, alternative firmware sources. But, boy.

Anyway, I thank Brian for bringing that Cisco research from their Talos group to our attention because it's exactly typical. And this is tip of the iceberg. It's just the problem is that everybody wants to get in on this IoT phenomenon and have fun with it. And the problem is that security's not their focus. It's selling, putting things in bubble packs on J-hooks in stores and saying, oh, look, you know, stick this in your refrigerator, screw this into your light bulb socket and so forth. And look at all the fancy things you can do. But security will just be a second thought.

In the last week, something took me over to the BleepingComputer site. We've referred to these guys often. They're neat guys. They were the early go-to site when ransomware first struck, and they have some online forums, and they had some good static remediation pages that talked about blocking the ransomware, detecting it, what to do about it. It's where we pointed people who wanted to spend more time digging into this. Something, I don't remember what it was in the last week, took me there, and I got hit with a page-covering pop-up that I didn't mind because it wasn't an ad. It was a plea. They apparently gave a negative review to something called SpyHunter, and in retaliation the SpyHunter people have sued them. So what they're asking for is some financial assistance.

If anyone has been using BleepingComputer, likes the site, goes there, takes advantage of it, give them a few bucks to help them with their legal defense fees because this is a lawsuit, basically frivolous, trying to force them to bring down a well-meaning, non-

malicious, but probably accurate review. I mean, these guys know what they're talking about. And it's unfortunate that, whatever SpyHunter is, it didn't get a glowing review; but SpyHunter ought to spend their effort fixing their product, rather than suing people who talked ill of them. So anyway, I wanted to give our viewers a heads-up in case anyone hadn't been there recently, but wanted to give BleepingComputer a hand because it seems like the right thing to do.

Boy, a new attack vector for DDoS amplification attacks. We've talked about DNS, how the problem with DNS and with NTP, the Network Time Protocol, is they're both meant to be very lightweight protocols, so they operate over UDP. You cannot spoof a source address if you want to use a TCP service because part of what TCP does, the TCP protocol, is that the so-called three-way handshake, you send a SYN packet, short for synchronize, to the other guy. They send back a SYN ACK, which is basically a compound packet, their own SYN, their own synchronization, and an acknowledgment of the receipt of your SYN. Then you finally send an ACK which acknowledges the receipt of their SYN ACK.

So that three-way handshake establishes the connection. But that also inherently validates the IP, that is, both ends know that the IPs that they have are accurate because they've just exchanged packets back and forth. So you can't spoof a TCP connection. If all you wanted in the old days was a SYN flood, you could send SYN packets off to a bunch of TCP servers and have them all send SYN ACKs back to the target of the SYN packet in a TCP reflection attack. That could be done.

But UDP is different. The idea is it wants to be just a query and a response, very lightweight, not all this set things up and get things synchronized and count the bytes and all that. No. It's just here's my question; let me have your answer. Perfect for DNS. Here's a domain name. Give me the IP. Bing, bing, done. Same thing for Network Time Protocol, very lightweight protocol. Well, there's another lightweight protocol. Nobody even thought that this was going to be a problem because there is no good reason for it to ever be public facing.

Leo: Which of course means it is.

Steve: Which of course means, right, right, Leo. And I don't know who is - and I'm curious. If some listeners want to find out, I'd like to know. Why are there 599,600, just shy of 600,000? It's got to be routers. And it's crazy. The protocol is known as TFTP, Trivial File Transfer Protocol.

Leo: Not SFTP. Not FTP. TFTP.

Steve: Right. Now, FTP is a File Transport Protocol that we're all familiar with, or at least old-school guys are. FTP is a TCP-carried protocol, meaning in TCP fashion you establish your communications, and then you have a nice command set that allows you to say things like "change to this directory," "list the contents of this directory," and so forth. Oh, and you have to authenticate. You have to give it a username and password in order to say this is who I am, in order to get access, to log into the FTP server.

But there's a class of need which doesn't, by definition, require all that. And, for example, it's when you have an old-school router. I've got a Cisco router, for example. And I did, I used TFTP, Trivial File Transfer Protocol, to update its firmware. I bought the

thing off of eBay, it was in good shape, but it had - it may not even have had firmware, or it didn't have what I wanted. So I found the right IOS, that's what Cisco's, its Internet Operating System (IOS) is called. And I got a simple TFTP client for Windows and hooked them up by serial cable, I think, or I guess it must have been Ethernet.

But the point is that it is UDP. It is a simple packet-and-reply protocol. There is no authentication in Trivial File Transfer Protocol, which again is how can there be 600,000 of these exposed to the Internet when there's no authentication available? It's like, who's got their TFTP servers, 600,000 of them? And they've of course been found by Shodan, and they're all there, just saying, "Hi. Ask, and ye shall receive."

So the Trivial File Transfer Protocol is exposed; 600,000 of them are out there. And they can create, it is estimated, and in a research paper which is behind a paywall - it's ScienceDirect.com. So I was unable - and I didn't want to give them 40 bucks for it. It wasn't that important. But the researchers developed a 60, six zero, factor bandwidth amplification with these TFTP servers, making it...

Leo: Sixty is not that bad, though; right? Or is it?

Steve: Well, it's up there. That's up there with the best of them.

Leo: Is it? Oh.

Steve: So, yeah, that's a substantial amplification. It uses well-known port 69. So somebody, probably Shodan, scanned all of the IPv4 space, sending UDP TFTP query packets to port 69, and got 599,600 replies. Meaning that there are - they have to be routers. I can't imagine - maybe they're light bulbs. Who knows? But something is out there saying, would you like to update my firmware? I'd be happy to have some new firmware, if you've got any ideas. Oh.

Leo: I mean, I guess you'd have to figure out what it is before you could...

Steve: Yeah, exactly. And the protocol doesn't require authentication. There's no list, I don't think you can list files of a directory. It is very lean. It was deliberately designed to be a minimal amount of code that you might stick in a ROM, for example, so that you could, you know, just enough, it's sort of like a bootstrap, just enough to then bootstrap the firmware so you could always count on the TFTP server being present. And then you would use it to then load the real firmware. You would then restart the device, and up it would come with the firmware that this little TFTP server had allowed you to feed into it.

So, wow. It'd be interesting to see, as this develops a little bit further, what those 600,000 devices are because, while they're useful for a factor of 60 amplification attack - because, again, it's UDP. So you'd simply spoof the source IP with a query to it, and it bounces its reply at your victim with a packet that is 60 times larger than the query you had to send out. So it must be a very small query and a very large response, obviously. I don't know the details, and I was unable...

Leo: That's what you want, though; right?

Steve: Well, it's not what we want.

Leo: All the amplification, yeah.

Steve: It's what the attackers are, like, drooling over. Although, frankly, if I were a bad guy, I'd be much more interested in what those 600,000 routers or whatever the things are, are.

Leo: Somebody attacked us the other day with a WordPress amplification DDoS. And it was a bunch of WordPress installs in Russia. Apparently it has to do with a WordPress pingback routine.

Steve: Wow.

Leo: Yeah. It wasn't a very good DDoS. I know I shouldn't laugh. But it wasn't a very good DDoS. But that just shows you there's all sorts of stuff out there.

Steve: Well, yeah. And that's the other thing, too. I mean, there are still protocol attacks. We've been talking about - or, I'm sorry, in protocol. We've been talking about sort of attacks at the plumbing level. But ultimately, a website wants to give you a page. That's what it's there for. And we've talked in days past. If a website has a very high cost to delivering a page, then you can hurt the server by asking for pages faster than it's able to deliver them so that, you know, that's not any kind of a filterable attack unless you have something, for example, caching those pages in front of your server. Which of course we talked about back when GRC was getting DDoSed. We talked about that was what CloudFlare was doing.

So last week I mentioned - oh, last week was Patch Tuesday, and it was happening just, of course, as Patch Tuesday does now, during the podcast. And I hadn't had a chance to look at it closely. One of the things that got a lot of people upset was that, in the guise of a security update, and this was KB3139929, that also included KB3146449. And the page for that 449 update reads: "This update adds functionality" - get that, functionality - "to Internet Explorer 11 on some computers that lets users learn about Windows 10 or start an upgrade to Windows 10."

Which of course we now know from the industry's massive backlash to this, I mean, people got a little over the top. I think it's just the people, as you said, Leo, at the top of the show, are just getting a little tired of Microsoft pushing Windows 10 as hard as they are. Woody Leonard, who's a well-known industry writer, he writes for InfoWorld now, he wrote: "Putting an ad generator inside a security patch crosses way over the line." Now, I wouldn't call this an ad generator. I mean, it's an ad probably only for Windows 10. But again, nobody...

Leo: I agree with Woody. I agree with Woody. It's way over the line. This is terrible.

Steve: Yeah. Nobody doesn't know that Microsoft wants them to have Windows 10.

Leo: But you can't, because it's a critical security update, you can't not get this GWX ad.

Steve: Right, right.

Leo: So annoying.

Steve: And I haven't mentioned this, but I should, and that is that I've given people the link, my little bit.ly shortcut, to what I have now proven is a bulletproof perfect solution. And that's the bit.ly/no-gwx. That simply expands to the Microsoft-sanctioned way of saying no.

Leo: Does it work?

Steve: Yes, absolutely.

Leo: Okay. So you don't have to get the GWX Control Panel anymore. Just use this.

Steve: Nope, none of that. None of that. This completely...

Leo: And they promise not to update on top of it and put it back.

Steve: Well, yes. But what I wanted to mention was my back was stiffened when I'm reading through the supported systems, and I see Windows 2008 R2 Server. It's like, whoa. Hold on. You mean....

Leo: Yeah, you don't do that to a server. You don't do that to a server.

Steve: Yes, they do.

Leo: Oh, that's terrible.

Steve: And that's my point. That is un-effing believable that Microsoft would push Windows 10 to my, to GRC's production server in any way.

Leo: That's horrible.

Steve: It's just like, I'm just speechless. And it's like, this is just like, well, okay. There it is. I'm speechless. Although, on the funnier side, to go from this dark black topic, I got a tweet from Chris, whose handle on Twitter is @selfuntitled. And I just - this just made me chuckle. He found a project on GitHub which is the anti-adblock killer. And so what he tweeted was so perfect. He said: "The ad block arms race escalates again, an ad block blocker blocker," which of course this is, an ad block blocker blocker. So this is the, for sites that are trying to block your ad blocker, now GitHub is offering you the ad block blocker blocker.

I've had a number of people ask me about setting up Windows 10 machines. I shared my experience fighting with those laptops where the Windows 10 came with Windows 7 SP1. Lord knows what happens if you just try Windows 7. It would really be out in the weeds. But even the Service Pack 1 is the final edition that Microsoft offers. And of course it's still, it's like 200-some security updates behind. But it's so old that they changed Windows Update after it in a way that causes it to get lost in the weeds.

So I created a new bit.ly link for people who want the basically Windows Update Update. And naturally it's bit.ly/wupup, Windows Update Update, or Windows Up Up, wupup. That will always take you to the Microsoft Knowledge Base article titled "How to update the Windows Update Agent to the latest version." And you could also just google "Windows Update Update," and it'll also take you to the same page. That is, that's how I got it from Microsoft. So, you know, there are lots of...

Leo: Did you see that a lot of people are complaining, and I've heard from a number of people, that Windows 10 upgrade is just installing without permission. A lot of people are starting to report this. Now, remember, we talked that they've moved it to the recommended update.

Steve: Oh, I - yes.

Leo: But we'd always said, and everybody had confirmed - and I'm going to have to ask Paul and Mary Jo about this - oh, no, but don't worry because, yes, it will download it. But before it completes it'll say, now, I have it. You want to upgrade; right?

Steve: I'm glad you mentioned it because I've run across those stories in the last few days. It's literally, it is rebooting their computer and taking them to Windows 7 without their permission.

Leo: Windows 10, yeah, yeah.

Steve: I'm sorry, to Windows 10 without - yeah.

Leo: And two people in our chatroom, I mean, our studio are saying it happened to them. Happened to you? Happened to them.

Steve: Oh. Oh.

Leo: That, you know, that's a bug. I don't think that that's what...

Steve: Leo, is it a bug, or is it a feature?

Leo: It's a bug. There's no way Microsoft wants that to happen.

Steve: Wow.

Leo: That's just bad.

Steve: Wow. Well, we know how badly, I mean, nothing could lead us to believe that they desperately don't want everyone to do that. I can't believe they're wanting to do that to my server. That sends chills down my spine.

Leo: No. Imagine if it upgraded without your permission.

Steve: Right. Who knows if my stuff's still going to work? Suddenly it just crashes, and it's gone, because it got Windows - what is it? Do they call it Windows 10 Server?

Leo: No, Windows Server 2016 or something like that, or 2015.

Steve: Yeah, 20 goodbye.

Leo: Goodbye.

Steve: Wow.

Leo: That is really - it is awful. It is really - that is just unacceptable.

Steve: So I did want to mention to people who have been trying, who have been receiving the Healthy Sleep Formula ingredients, that I'm getting - I'm beginning to get feedback from people who are reporting the end of years of insomnia. So it's batting a thousand. Some people are succeeding adding Taurine. For some people it is actually stimulating, and it hurts their sleep. Taurine itself is so good that they could easily just

consume that bottle they purchased in the mornings, and it would be good for them. I'm seeing a cessation of all of the ringing in my ears, the tinnitus that I've had. And it has...

Leo: What?

Steve: It has some of that reported effect.

Leo: I have bad tinnitus. So if that happens, I'll be really...

Steve: So you might just take...

Leo: I'll send you a bottle of your favorite burgundy right up.

Steve: Take extra in the morning, Leo, and then see if...

Leo: Taurine.

Steve: Taurine.

Leo: Taurine.

Steve: Yup.

Leo: Wow.

Steve: And I did hear you also during the preshow of The New Screen Savers. You had an unsolicited discussion about the Healthy Sleep Formula.

Leo: Yes.

Steve: And it's continuing to work for you.

Leo: Loving it. Now, but this is what worries me. I'm kind of, like, come to Daddy. Okay, come to Daddy, sleep formula, I'm ready to go to bed. It's like I love it.

Steve: I know.

Leo: You got me hooked.

Steve: I know. It's just you look forward to sleeping now.

Leo: I look forward to it.

Steve: Because it's just going to be a really nice night of sleep.

Leo: It's not the struggle. It's not the struggle it used to be.

Steve: Yeah.

Leo: We sound like we're doing an infomercial.

Steve: Well, the good news is it's free. And that's the point. The FDA has nothing to say about this because what they regulate is commerce.

Leo: Well, they've always been told that they can't regulate supplements, which was of course the supplement association of America or whatever getting a law changed. I think they should regulate supplements, but...

Steve: It would be, as long as - if it didn't increase the prices...

Leo: Well, it would. And it'd decrease the availability. It would cause all sorts of problems. So it's probably, you know, just - but consult your physician before you do this stuff.

Steve: I think it's Orrin Hatch. Is he still alive?

Leo: Yeah. Is that Oregon?

Steve: Yeah, I think Orrin has a huge amount of stock in the supplement industry, and so he blocks all the legislation that comes along.

Leo: Is that true? Wow.

Steve: So Orrin, you go, boy. You just keep blocking any attempt to...

Leo: Senator, Utah, yeah.

Steve: Yeah.

Leo: Oh, well.

Steve: But anyway, I did want to mention - you also mentioned during The New Screen Savers preshow that you were loving Season 4 of House of Cards.

Leo: Yes.

Steve: Now, I stopped after Season 2 because I'd heard it wasn't any...

Leo: Three wasn't very good. Three was disappointing.

Steve: Well, and I have to say I guess they wrote it for me because...

Leo: You have to watch Season 3 before you can watch Season 4. It's all...

Steve: Well, that's why I started watching Season 3. My point was I love Season 3.

Leo: Well, you'll love Season 4.

Steve: Well, okay. Because, you know, those first two seasons were like, you know, people were being asphyxiated with, you know. And it's like, okay, this just seems a little farfetched. Three seems pure political machination.

Leo: It is. Oh, you're going to love Season 4. Oh, man.

Steve: Neat, neat. And I agree with you. Robin Wright is just - she's just acting herself to a fabulous...

Leo: Wait'll you see Season 4. She's in many ways now - that was probably what was wrong, in a way, for Season 1 and 2 was you couldn't really understand what's her role in all of this.

Steve: Yeah. And their relationship was odd, too.

Leo: Yeah.

Steve: You know, they would sit...

Leo: Yeah. You get a lot more insight into it now.

Steve: They would sort of sit in the windowsill and share a cigarette. It's like, okay, well, but...

Leo: No, no, you know what? I'm now of the opinion that they had planned this whole arc all along, and this isn't like an afterthought. This explains a lot about Season 1.

Steve: Interesting. Oh, I can't wait. It's like...

Leo: I haven't finished it yet.

Steve: Okay, thank goodness for Healthy Sleep Formula because I need it. And I need to say that I continue to look forward every Monday night to "The Magicians" on Syfy. I just like it. Maybe they could pick up the pace a little bit. But I really am enjoying it, so for what it's worth. And I think Question 7 there was a comment about someone in his - I think we might have had two questions about hard drives. So I thought I'll wait to talk about SpinRite till them.

Leo: A Q&A. Feels like, again, we went a long stretch without one. I don't know why. Maybe we did. I can't remember. But I've got them for you. Did we do questions last...

Steve: We've been covering so much news, I think. There's just so much going on right now.

Leo: So Dave Redekop was here a couple of weeks ago during the Security Now! show, watching the show. Great guy. He was a sponsor years ago with Nerds On Site.

Steve: Yup.

Leo: We've always thought the world of him and his business. And he said, he wrote: I thought you and your audience would like this tip, StealthyLinks.org. Sometimes, he says, it's useful to link to a bad site from a high-value site. But we don't want search engines to assume that the linked-to site is any good.

Steve: You know what he means?

Leo: Yeah, you don't give them any Google juice.

Steve: Exactly.

Leo: Now, Google has a way to do this, I should point out, but you have to have access to the HTML. You just rel="nofollow" and then it gets no Google juice. And that's what this guy's doing. So the site allows you to create a link that will bounce the user's browser through the site. See, this is the part I don't like. But it's different from a regular URL shortener in that it's purposely obfuscating the source and eliminating trackers en route. All the best from green Canada. Thanks for the awesome podcast, as always. @DRtheNerd.

So, yeah, very simple. In fact, a lot of software, like WordPress and others, will allow you to say, "Make my links nofollow." But you can also do it by hand, add nofollow in there. Then it goes straight to the site, but they get no Google juice from the site.

Steve: Could you post a nofollow link on a social media site? I mean, and like add the nofollow? You may not be able to control that.

Leo: Oh, I see what you're saying, yeah. Okay. Okay. I'm sure bit.ly and others would let you do that. That's an interesting question. I'll have to google that around. That's a good point, right, you're posting it into Twitter, but you don't want to give them any Google juice.

Steve: Well, or you just don't have, like, all I can think of - sanitize is the word. I was thinking sanctify. No, that's not what I mean. Sanitize.

Leo: Twitter does no - Twitter, I'm pretty sure, does rel="nofollow." I'm going to have to look this up. I'm pretty sure Twitter does rel="nofollow."

Steve: So, I mean, they add...

Leo: They add it.

Steve: They preemptively add that, ah.

Leo: Yeah. This was a big deal because - and by the way, you're not getting much social Google juice when you put something on Twitter or Facebook.

Steve: And I was going to say, so those are examples. But just in general, like putting a link in a blog posting reply.

Leo: That's for sure what you would want to do in a blog. On the other hand, if you want to promote the site you're linking to, which mostly you do because you like it, don't put rel="nofollow" in there. Yeah, Facebook and Twitter have added nofollow to the links.

Steve: Oh, nice.

Leo: So you don't have to worry about Facebook and Twitter. They're not getting any.

Steve: So it's neutral.

Leo: It's neutral, yeah. Oh, I'm sorry, more questions. Thank you, David. Let's get back to the Q&A.

Steve: And good luck with this location in Wisconsin. When I transcribed this, I thought, okay, Leo, you are - maybe you actually know where this is.

Leo: Everybody knows it's Oconomowoc.

Steve: Wow.

Leo: Michael in Oconomowoc, Wisconsin notes another crypto deprecation. I'm good at pronunciations, pronouncing things. As hard as I try to listen to every episode - I might be saying it wrong, by the way. I have no idea. I just, if you say it with confidence...

Steve: That's right.

Leo: Oh, you mean Oconomowoc? Yeah. For all I know it's pronounced Wooster. I don't know. As hard as I try to listen to every episode, a few slip by, and you may have already addresses this. I'm reading it exactly as written, [sic]. Addresses this. I thought you and the community would like to know that the Payment Card Industry, or PCI, standards group has moved their target for TLS v1.0 deprecation - speak of the devil - from June 2016 to June 2018: "The Payment Card Industry Security Standard" - which is known as PCI SSC, oh, I'm sorry - "the Security Standards Council is extending the migration completion date to 30 June 2018 for transitioning from SSL and TLS 1.0 to a secure version of TLS," that is, 1.1 or higher.

Steve: And I thought this was really interesting because it's the first that I had heard that we already had plans afoot to abandon TLS 1.0. It's like, wait, we just got here, and we're leaving already, because of course we have 1.1 and 1.2. And if we're learning anything, it's that it is just so difficult to get people to get off of what's working. Even

though there's increasing evidence that it is insecure, people just still want to use what they've got. There was a dialogue with Ryan Sleevi, Google's Ryan Sleevi, over in the CA/Browser Forum, with someone still arguing about SHA-1 and not being happy about being forced off of it for some reason. It's like, boy, you know, that ship has sailed. You just have to - in fact, speaking of the devil, I'm now running SP3 of XP.

Leo: What?

Steve: I bit my own bullet over the weekend.

Leo: What? Welcome to 1992, Steve.

Steve: Exactly. I know. Ever since I updated GRC's certs on New Year's Eve, I have not been able to bring up GRC, my own site, on my own machine, under IE. I use IE because it's just so lightweight, and it's just sort of there. It runs fine under Firefox, and I don't even remember now Chrome because - I think it did on Chrome because remember Chrome uses some of the native platforms' security stuff. But finally what happened was that SQRL, my SQRL client does use the built-in crypto in Windows. And SQRL wouldn't work. So I thought, well, that can't stand.

So long ago I had tried to install SP3, and it just sort of didn't like my system. And I always thought that was probably because I had turned off a bunch of services, you know me, trying to run as lean as I can. I went through the service list and said, what is this? I don't need that. Turned it off and disabled it and disabled it and disabled it. And as a consequence I think that's always why SP3 wasn't happy, but everything was working fine, and I left it that way. But SP3 was when XP did add awareness of SHA-256 certs. So now SQRL is running again, and I can access my own site from my desktop.

Leo: Wow. Nice.

Steve: So miracles never cease. However, now we know that even TLS 1.0 is destined to be removed. But clearly summer, the middle of this year, June 2016, nah, that was never going to happen. I don't know what the back story is behind them punting two years. But that seems reasonable. I mean, there are no known problems with v1.0, or we wouldn't be using it now. As far as we know, it's just fine. We got away from SSL.

Leo: Well, doesn't it tie into POODLE? Can't you use a downgrade to get to SSL3 and POODLE?

Steve: There is, well, next week Father Robert and I are going to talk about DROWN, D-R-O-W-N, which is an interesting hack that does use access to an older version of SSL. But here we're just talking about TLS v1.0, although that is almost the same as SSLv3.0. So I'm glad that they're moving forward. And I'm glad that they're giving people enough time. So, yeah, it just, I mean, it wasn't going to happen by summer. We're just - we've barely recovered from giving up SHA-1. That was enough of a workout for everyone. And I'm finally back online. So it's like, okay, good, let's leave things alone for a while.

Leo: A wonderful thing. Let's see. Wait minute, let me go back to the questions.

Steve: Scott.

Leo: Scott in Columbia, Maryland. He's wondering a little bit about encryption, the cloud, and security: Steve, I have a question regarding exactly how secure it is to store encrypted information in the cloud. Well, you've got to use [P]. I have a file - that's me, by the way, not Scott. Scott wouldn't say something like that. I have a file, obviously containing sensitive information, which is currently kept in a VeraCrypt volume on a thumb drive. The VeraCrypt volume was set up using the AES-Twofish-Serpent algorithm, which, if I am reading the VeraCrypt docs correctly, means the container is thrice encrypted.

Steve: Oh, yes.

Leo: And that's better than twice encrypted, which is almost as good as single encryption. The password for the VeraCrypt container is a very long random string obtained from your Perfect Passwords page. Good. The password is of course stored in LastPass. You didn't think I could possibly remember that, did you? My question is would it be safe, in your opinion, to store a copy of this VeraCrypt container on a cloud service like OneDrive or Dropbox for backup of the very sensitive data - taxes, for example? But just parenthetically, let me ask you, Steve, is it better to encrypt something three times than it is to encrypt it once?

Steve: Okay. So what struck me was the overkill of this. Now, the argument - so what he's done is his data is encrypted with Serpent, which is a good cipher. Then with Twofish, which is a good cipher. Then that is encrypted with AES, which is the industry standard cipher. The argument for multiple encryption is - and, now, maybe they each have different keys. He didn't specify, or whether they are using the same key. If they have different keys, then you end up with a key three times as long, which is like crazy long. But the argument for multiple encryption is, if a flaw were ever to be found in one of those, then your data is still secure. So that's the argument. However, standing back from this, he's got this thing triply encrypted and wants to know if it's secure to store it in the cloud. My take was the weakest link was LastPass.

Leo: Right.

Steve: Which is not to say that LastPass is weak. But the extent to which he's gone to protect this thing, which must be the crown jewels of existence, you know, and a crazy long, absolutely high-entropy password that he got from GRC's Perfect Passwords Page. You do not want to put it into LastPass because, I mean, I love LastPass. I use LastPass. I recommend LastPass. But in terms of the weakest link, that's the weakest link because it's...

Leo: And why is that?

Steve: Well, because it's a browser. It's running JavaScript. You're depending upon nothing in your system ever. You're depending upon your system never getting infected with something that had a way of worming its way into LastPass. I have no way of knowing if there is any way. But just from an analysis at 10,000 feet, the defect in his exquisitely over-compensatory encryption is he stored it in LastPass.

Leo: Yes.

Steve: Because, which is not to say it's not secure. But of everything he's done, that's the weakness. And so I would, by all means, I know that my pseudorandom generator is very secure, pseudorandom. No one has GRC's keys. That password went from me to him. No one could get it. I would use it, but then I would store it on a thumb drive and put it somewhere else. And print it out because he doesn't have to remember it. It's not going to be that long. But if he absolutely had to type it in again, he could. But the thumb drive is more convenient. The printout is absolutely bulletproof against - prevents against loss. But don't store it on your local computer in any form, even in LastPass, because that's the point of greatest vulnerability. So Scott asks, "Is it safe?" Oh, my god, yes. Triple encryption with different algorithms? That's very safe.

Leo: Jay W., Washington State, wonders about Jupyter in the web browser: I'm a chemical researcher and instructor. About half a year ago I began using a tool called Jupyter - it's J-U-P-Y-T-E-R, Notebook, you can find it at Jupyter.org - to assist my research and teaching, and I find it immensely useful. It allows me to generate documents containing code, documentation, graphics, and other elements to organize and share my work with students and other researchers.

However, I am slightly uncomfortable with the fact that the notebook runs in the web browser and connects to a server created by the Python kernel. If browser inter-tab sandboxing and my firewall prevent malicious access to my notebook and server, there should be no problem. However, the fact that I'm writing and executing Python code from the browser makes me uneasy.

As an extra precaution, I'm thinking it would be nice to have a separate browser that does not connect to the Internet, that I can dedicate to Jupyter. Would this be an effective precaution, and how might this be accomplished? I don't see an easy way to configure a Mac firewall to block a specific browser. I'm also curious about what it means that the browser is connecting to a server on my system? Is this something I should be concerned about? Thanks for the great podcast. Jay W. from Washington. I'm going to have to take a look at this Jupyter thing. I'm not sure what it is.

Steve: You should, Leo. It's a really - I was unfamiliar with it before he mentioned it. And it's a very interesting system. It looks like, it says it supports over 40 different languages. The server is written in Python. And it looks like you're actually able to author code in the browser that it then runs, so in order to, like, create animations and physical simulations and all kinds of stuff.

Leo: Wow, this is really cool.

Steve: It really is.

Leo: So what are the security concerns here?

Steve: So, okay. So what they've done is interesting. They wanted a cross-platform solution. So they implemented their system as a server running in the local system.

Leo: Localhost.

Steve: A localhost, exactly, which you connect to with your browser. Now, it happens that I've just come out of several months of work with the SQRL group, looking at the idea of the SQRL client in a user's machine creating a localhost server so that the browser could receive the session permission, essentially the session tag, directly from the client, rather than looping through the Internet. Because that would, if it were feasible, would provide absolute man-in-the-middle protection.

The problem is Microsoft is deprecating the use of localhost servers. They're threatening not to allow it in their Edge browser. In fact, for a while it was going to be turned off. Then before they shipped they turned it back on. But they really seem to be making noise that they don't want to allow this in the future. So unfortunately we developed the technology, we proved it, and we decided, okay, the proper place to do this would be with a plugin. A plugin would have the same capability of establishing its own communication. We were just trying to do it in a browser-neutral, plugin-free way.

But the point is there's nothing insecure about this kind of communication. I also saw something in the last couple weeks where some other security product was using the same thing. That is, they were setting up a server, listening on 127.0.0.1, the so-called localhost IP, and some other component of their system was talking to that server through the localhost. Which it makes people nervous because they sort of think of it as the Internet. They think of it as IPs. And we know of all the scary things that can happen between IPs.

But the way Windows and other operating systems implement this is it's basically interprocess communication. It is, oh, I know what it was, it was something that was going around where there was this communication between processes that was not encrypted. And it was like, you know, shouldn't it be encrypted? It's like, well, it actually all occurs in the kernel. And while, yes, you could technically intercept, if you modified the kernel and intercepted kernel-level things, then you could intercept the communication between that localhost server and the client that is accessing it. If you have that capability of being in the kernel, then all bets are off anyway. You could do anything you want to, if you're able to - if you have that level of control over what's happening in the kernel.

So I don't see anything wrong with this notion of something is running a server in the local machine, and then the browser, you probably put into it, you know, http://127.0.0.1, hit Enter. And maybe colon something, like it may not be running on...

Leo: Yeah, it has a port, 8888. It's port 8888.

Steve: Ah, right. So that's a common fallback from the normal HTTP port 80. So :8888. And then it's going to, like, bing, now you're looking at a page generated by that local

server. Now, his concern is what about intertab isolation? And I agree with him that we are depending upon the browser not making any mistakes and keeping code in one page from a different domain, like some malicious.org, separate from 127.0.0.1. I wasn't sure when I read this whether the same-origin policy did incorporate a port, and it does. So because it if didn't, then you could - some malicious person could set up 127.0.0.1 on a different port, and then have the same 127.0.0.1 IP and then break intertab isolation.

But same-origin policy, as we would hope, these guys were on the ball. They include the port. So what that does is it prevents that attack. And same-origin policy, which again, remember, it is the first thing browsers do is prevent queries across origins because that's the only thing that keeps tabs safe from each other and keeps them from playing any games with each other's cookies and so forth. So which is why that Chromodo browser, which didn't enforce same-origin policy at all, just had the Google guys breathless with disbelief. It's like, you took our Chromium browser. First of all, you named it Chromodo, thank you not. And then you ruined its security.

Yeah, anyway, so I don't think Jay has anything to worry about. This is a very cool application. And it's sort of nifty that it - so the server being in Python means that it's portable, so there's probably a version for Mac and Linux and Windows and anything. And then you just use any browser to communicate with it, kind of a cool way of doing a very high-end, sophisticated UI just using the browser, rather than like a custom app.

Leo: Yeah. In fact, thank you for tipping me off. I'm downloading it now, installing it on my server. It's a great idea.

Steve: It's neat-looking, yeah, yeah.

Leo: Yeah, because you don't have to have it on a local server. You could have it on a regular server.

Steve: You could.

Leo: And then you'd have your own web-based Python notebook or Rust or R or whatever else you wanted to use.

Steve: Right.

Leo: You use it? You're nodding. But, yeah, it's cool. Yeah, it's a really neat idea. Python folks are really amazing. Python is such a great community.

Steve: I had a very good friend who I've known since my early 20s, one of the sharpest developers I've known, who said to me a couple years ago - and he's been coding straight for 40 years. He said, "I have never been as productive as I have been in Python."

Leo: Yeah. Oh, it's amazing, yeah.

Steve: It's just, you know...

Leo: Lovely. Lovely code.

Steve: Lord knows there's a bazillion languages out there. And when someone now says to me, "What should I learn first?" I say take a look at Python.

Leo: Yeah.

Steve: The nice thing is that now the whole ecosystem is so huge and mature that you can get all the help you could ask for. Anything you want to know, the 'Net has it. Lots of tutorials. Lots of references and information onsite or online. And great O'Reilly books, too.

Leo: And they use a Python called IPython, which is an interactive Python, designed, I think, for Jupyter. Scientists love Python because the science libraries, the math libraries are so phenomenal. Yeah, great language. Let's move on. Question number whatever here. Wait a minute, let me find it.

Steve: Five.

Leo: You know why, because I started installing Jupyter, and I lost the questions here.

Steve: Focus.

Leo: Question 5 from Jim Murchison. Focus, I know. Mountain Ranch, California. He found a frightening SSL cert in his system store: Being one who checks his SSL certs every so often, I recently found an "all purposes" cert in my certificate store from "Applian Replay Media Catcher 6" that was essentially - now, by the way, Replay Media, I know this program. It's a good, useful program. You can use it to record audio and video that you're streaming on the 'Net that isn't intended to be recorded. And I suspect that's why they did it this way.

Steve: Uh-huh.

Leo: It's essentially a man in the middle sitting on my own machine, much like the certs from antivirus companies. Do you know of any tools that monitor cert stores and report when changes are made?

Steve: So, if I weren't so far behind, I would immediately whip out a utility. I'm not going to.

Leo: That's a great idea, yeah.

Steve: It is a great idea.

Leo: Someone should write that.

Steve: And that's why I put this in this Q&A. Remember that Mark Russinovich just recently updated the Sysinternals tool so that, with a command line, it'll do that. Somebody could write a little Python front end to the Sysinternals tool that is invoked by the scheduler, or maybe just runs in the background and checks every day. If anyone does, make sure you send me a note, and I will make you famous. I will tell everybody about it because that ought to exist.

I think the idea of something keeping an eye on our root store for anything that comes in that, well, any changes, first of all. You could certainly do it that way. So you manually audit it yourself; then say, okay, I want to be informed of anything that changes. Or do what Mark did, which is obtain the official store from Microsoft and do a comparison. So you would be alerted to any additions to the differences. Because, for example, again, I'm sure that Applian Replay Media Catcher 6 has got the best of intentions. But they also have an all-purposes cert, which makes it - it's a god certificate. It can do anything it wants to. And it's in your certificate store.

So what that presumably means is that your system is now trusting anything signed by it, which means anyone could reverse-engineer that application and produce certificates that your computer would trust, even when it should not. And that's just not the way it's supposed to work. This is an abuse of, unfortunately, as you said, Leo, this may be the only way they had of doing what they needed to. But this is an abuse of the root store. We just don't want random companies dropping all-purposes certificates in that our system will then trust. This is not a good precedent. So again, if anyone out there is a coder, wants to come up with something that manages our cert stores, let me know if you've done that, and we will absolutely let everybody else know.

Leo: Couple of them that you might be interested in. Strengths mentioned that there's a Firefox add-on called Certificate Patrol. I think we've mentioned that.

Steve: The problem with that is that Firefox has its own private store.

Leo: It's not the systems store, it's patrolling the Firefox store.

Steve: And Certificate Patrol's been around forever, by the way. It's a nice add-on, but it's not going to do the job for us.

Leo: Will Sigcheck, the Sysinternals tool, will that check just the browser store? Or will it go...

Steve: It checks the system store. And that's certainly what these guys, the reason they put it in the system store is they're trying to intercept the system's security.

Leo: Right. Yeah, and I think that what you're going to see is there are, if you install something, is that there are a lot of tools that need to do this as a way of doing something unusual, as in the case of Replay. It's saving a stream that's coming in from the browser that is marked not to save.

Steve: Right.

Leo: And so they have to kind of intercept it and download it. And I don't know if it's the only way to do it, but I can understand why they did it that way. CJ in Indiana wonders about - and if you took that cert out, it would stop working, we should point out.

Steve: Yes. And I should just mention it's not that I'm taking any absolutist position. It's just information.

Leo: Good to know.

Steve: You know, for example, this should not have come as a surprise to Jim. He should have been informed by this thing and gotten its permission before doing it.

Leo: Typical in the Mac world, though, because they don't want - oh, incidentally, we're going to put in a special certificate so we can intercept traffic so that...

Steve: I know. So, like, why is there a dialog here? Okay, okay, okay, okay. You know...

Leo: Yeah, yeah, okay, yeah, that's fine, I don't care.

Steve: Get rid of these dialogs.

Leo: Whatever you're saying, just stop saying it and let me use the program.

Steve: Yup. Yup.

Leo: CJ in Indiana wonders about the necessity of user agents in browser headers: Steve, you've talked a lot about user agents being disclosed in the browser's request headers, and how this can be used for fingerprinting and possibly to identify targets. How necessary is this information in the browser header? I feel like we'd be better off if the browsers stopped advertising nearly the amount of information they do to every website we visit. Do you have recommendations on how to limit the information in the browser header? Thanks. CJ.

Steve: So, okay.

Leo: Good reason for them to do this.

Steve: Yes. There's two parts. First of all, I don't - okay, yeah. First of all, I don't really think it's a problem because, I mean, technically, maybe someone could fingerprint you based on the unique or not necessarily unique nature of the headers in your browser. But it seems unlikely. I mean, even Panopticlick, which is all about auditing how common they are, when you go to Panopticlick and have them look at your browser, as it does, it'll say, oh, you're - we've seen 33,426 different fingerprints, and you're number 32,197. So the point is you're not unique, Snowflake. But you're one of a group.

So, I mean, so first of all, I wouldn't get too worried about the fact that our browsers have all this junk that they're sending out as request headers. I do think some of it is superfluous. Once upon a time, before we were as standards-based as we are, there were servers that would look at the user agent and send something different, depending upon which user agent it was. And I remember our screen resolution used to be sent, but that got deprecated some time ago because it just became irrelevant, or because JavaScript was now able to read that same information and do whatever it wanted to do. So the first part is I don't think it's that big an issue.

Second is there are privacy plugins, probably for any browser you've got. He didn't mention which browser he's using, so I didn't go do any digging. But I know for Chrome and for Firefox and probably for IE there are add-ons specifically for obfuscating. Some even randomly change the user agent, just so that every query you make comes out looking like something else. So if you wanted to play that kind of game, you definitely could find something for your browser that would strip things out that are not necessary, or just change them on the fly, just to throw off anybody who's trying to fingerprint your browser. But again, it's like, eh. I just don't think it's a huge thing to worry about.

Leo: And there are many other data points they can use to fingerprint you.

Steve: Exactly that.

Leo: You know, take one out, they'll find another one. Dan Eshleman in Washington, D.C., and probably listening right now from his secret lair in a five-sided building, has a question about SpinRite and SMART drive data. Steve and Leo, first of all, thanks for an awesome podcast. My coworkers and I listen live from the office every Tuesday. Hello, Dan and coworkers.

I have a question about SpinRite and SMART drive data. My parents' four-year-old iMac is a bit slower than I would expect it to be, so I thought maybe there's some bad sectors on the drive. I tried a utility called "smartmontools" to take a look at the SMART drive data. And from what I can tell - I'm sorry, I shouldn't laugh. And from what I can - it's just SMART drive is so dumb. And from what I can tell, the drive is operating normally. Are there any indicators in the SMART data that mean it's time to SpinRite the drive?

Steve: Okay. So a couple things. I have a lot of experience with SMART because SpinRite v6, it was the first version to dynamically monitor the drive's SMART data. And what I saw was that the health indication that SMART surfaces is relatively short-term. That is, if a drive is having problems, running SpinRite on it will depress the health...

Leo: I'm so depressed. Oh, sorry.

Steve: ...indication because the drive is having problems reading data. But a few hours later, that SMART health indication will have recovered. Meaning the drive has noticed, oh, a lot of time has gone by, and I don't seem to be having any more problems, so I'm going to report that I'm feeling better now. Which is one of the coolest things about the way SpinRite interacts with the SMART data because that SMART health should not be depressible by SpinRite. So if it is, then it's because the act of asking the drive to do some work is reporting I'm having problems here. But unless you ask it to do work, it isn't having problems because, after all, work for the drive means reading and writing stuff.

So, first of all, looking at static SMART data, it can in some cases show you a problem. It certainly could be that something like the relocated sectors parameter is, like, really, really low. But here's the second part is how do you know it's low for that drive? Because there are, unfortunately, no standards in the SMART industry, or in the SMART spec. Back when this was created, under mandate by Compaq, that had a lot of market clout back then, they forced the industry to implement SMART. The industry didn't want to. So it did the worst job it possibly could in creating SMART.

Leo: That's the real...

Steve: They considered it their tech...

Leo: That's the real story right there.

Steve: Yes.

Leo: They didn't want to do it.

Steve: Yeah. They created the weakest, worst, most useless thing they could that would satisfy Compaq, who wanted some means of knowing what was going on in the drive, yet

wouldn't disclose what was really going on in the drive. So, A, you can use SpinRite. And after it's been going for a while, look at the SMART page. And what I did is I take the, deliberately, understanding the way it works, I use the starting SMART health parameters of whatever is published, and any decrease in those that occur I show as red because there shouldn't be any. But if SpinRite does manage to push some of the health down, it shows it very clearly as red appearing in the bar that tells you, okay, this drive is saying "I'm working harder than I should."

But absent that, the only thing you can do is compare it to an identical drive. So if you've got smartmontools, but don't have SpinRite or can't for whatever reason run SpinRite, but you do have an identical system with the same drive, compare the smartmontools output of both. If they're the same, then it's probably the case that you're not learning anything. But if they're very different, if one of them has much lower health parameters, then that's an indication absent any dynamic work that you're asking from the drive, that this drive is actually in worse health than the one you've compared it to.

But again, it's got to be the same make and model of drive because there's no standard. It's only the same make and model of two drives will be the same, but not any other manufacturer, and not even other drive models from the same manufacturer. It's just completely - it's almost worthless except that SpinRite, because it uses the SMART delta, that is, the change in SMART data, that ends up being highly diagnostic.

Leo: Question 8, Rik Sagar in San Jose, California wonders why he doesn't need "Let's Encrypt" on his blog. He's talking about the Let's Encrypt certs. Steve, I can't believe you said, quote - he's obviously pasting from Elaine's transcript: "... well, yeah, I mean, it's a blog, and nobody needs security there."

Well, you're right, I don't need to protect the words of my blog. But as you have taught me in the past, and as also mentioned in this week's show, we need HTTPS to protect people visiting the blog from a person in the middle, meddling. Example: Verizon supercookie injection of malicious script tags into the HTML.

Keep up the good work. Glad I haven't heard anything about your DDoS this week. @riksagar on the Twitter.

Steve: So this was interesting. That's apparently a quote. So let me put it in context, which it needs. What I was saying was, historically, if the barrier was high, that is, you had to purchase a multi-hundred-dollar every couple years certificate, and figure out how to install it, maybe get a hosting provider to do it. And in some cases hosting providers, before we had SNI (Server Name Indication) in our TLS connections, some hosting providers that were hosting many different websites on a single IP were unable to disambiguate certificates, so you couldn't. Using SSL wasn't an option back then.

But my point was, for all of those reasons, if the barrier to using it was high on something like a blog that didn't have a high security need, it didn't happen. And it is happening now because the barrier to using HTTPS has just disappeared, thanks to Let's Encrypt. So it wasn't that I was saying you don't need it, it's that, historically, the sites that didn't need it that much didn't bother. And now everybody can to the point where I'm just like, no excuse not to at this point.

Leo: Although there is a not insignificant barrier in some cases because you have to

install the certs. Getting a cert for free is nice, but you still have to modify, I mean, putting...

Steve: No. No, Let's Encrypt does that. Part of what...

Leo: Oh, you don't have to do that, oh.

Steve: Don't even have to do that.

Leo: Right, you run that little script, right, yeah.

Steve: Part of what Let's Encrypt does is just magic. It edits your config files for you. You just sort of say, go, baby, and it says, I'm already done. And you're just like, what? And then you say, you know, log on with HTTPS, and it does. Yeah, I mean, they just - I think what we're going to see is it will just be built in, in the future.

Leo: I'm going to try it on my Minecraft server.

Steve: Cool, do.

Leo: Yeah. I mean, the problem is it's not Apache. It's not an HTTP server.

Steve: Oh. What is it?

Leo: It's a Minecraft server.

Steve: So Minecraft actually boots?

Leo: No, it's running on a Mac. And then you run the server software, just as you would if you were running Apache. Let me think about this. Could you do it at the DNS server level? No, you have to do it on the server server.

Steve: Yeah.

Leo: See, that's the thing is not everybody's running Apache or IIS. I don't know. What does Let's Encrypt work with? Probably just Apache; right?

Steve: And Nginx.

Leo: And Nginx. Okay, well, that covers a lot of the waterfront.

Steve: It really does.

Leo: Yeah, yeah. That's cool. No, probably there's no HTTPS for Minecraft servers. How about for Jupyter pages?

John in Montreal argued that Shamir, the S in RSA, well, and Rivest and Adelman, the R and the A, did not invent RSA. Or I think he's talking public key crypto; right? This distinction goes to some British cryptologists who came up with the idea during World War II. It was only declassified and made public well after RSA published their idea.

Steve: You know, I just sort of wanted to take this moment to comment that it is not true that the RSA guys did not invent RSA or public key crypto. They did. They just weren't first. But they did it without any knowledge...

Leo: Because it was classified.

Steve: ...of anyone else having done it. And so, I mean, this is an interesting area that we've run into because I sometimes hear people saying, oh, well, they didn't invent that because somebody else did. Or what often happens is two groups are doing this stuff at the same time, with no cross-communication, and they simultaneously invent, and then one publishes before the other. And to my way of thinking, that takes nothing away from the other group that independently invented something. So the fact that some cryptographers in Britain may have come up with public key crypto in World War II, but it was always kept secret, in no way diminishes the fact that the RSA guys independently came up with the idea. I mean, you know, there are lots of smart people in the world.

Leo: And of course - okay, I shouldn't say this. But in hindsight it's pretty obvious, but only in hindsight, right, because, I mean...

Steve: That's the best invention.

Leo: Yeah, right? Like, well, that makes sense.

Steve: Oh, my god. Velcro?

Leo: Right. In hindsight it's pretty obvious, yeah.

Steve: And a zipper? You know, you could spend a lot of time staring at a zipper, trying to figure out how that sucker works.

Leo: What's interesting about it, a lot of inventions would be hard to describe. But I think you can describe public key crypto very straightforwardly, in a sentence, in a prose sentence.

Steve: And I think that's what makes it so beautiful.

Leo: Yeah, it's elegant as hell.

Steve: Yes.

Leo: The idea...

Steve: You know, it's one of those head slappers.

Leo: Yeah, it's like, oh.

Steve: It's like, oh. Why didn't I think of that?

Leo: And it does make sense that anybody in crypto, going back for centuries almost, would be trying to find a solution to the symmetric key problem, which is you have to have a key exchange for a symmetric key to work. I have to somehow get you the key so that I can send you a coded message.

Steve: Right.

Leo: And that's a very big weakness because they're going to capture the courier with the key, and then they can read everything. So the idea that you could exchange messages with a publicly available encryption key...

Steve: Smoke signals, up in the air. Everybody can see it.

Leo: Anybody can see it. But without the private key, it's useless. It's brilliant. And you can see how the Nazis would have loved that instead of Enigma, which was by the way a symmetric key system.

Steve: Right.

Leo: And that's, in fact, you know, the fact that it was symmetric was how it was decoded, was that Turing knew that these signals were going out every day, these

radio signals with the key. And if you could - so his task was to figure out what the key was, you know, how to decrypt the key that's being passed out. And once you get that, then of course every day you get a new key, and it's a symmetric key, so you can decrypt all the transmissions.

Steve: Well, in the case of Enigma, though, they did have...

Leo: They had the machine.

Steve: ...the master code book.

Leo: Oh, they had the code book and the machine, yeah.

Steve: Yeah. And the book was what was captured.

Leo: And that was what you would use to decrypt the daily message, yeah.

Steve: Exactly. Exactly.

Leo: But you had to find the right page.

Steve: Because it had the machine settings for that day, right, right.

Leo: Oh, I love it. I love crypto.

Steve: And I just saw the movie again, by the way. It's just "The Imitation Game" is a great piece of work.

Leo: Great movie, although, again, historically, really a disservice in many ways to Alan Turing.

Steve: Unfortunately, yeah.

Leo: At least it popularized that whole notion. You know, it's interesting how encryption and crypto has broken into the public consciousness in some ways. And there's that great book, "The Codebreakers" by Kahn, which everybody should read.

Steve: Yup.

Leo: If you're interested in this, that and Steven Levy has a great book on crypto, I think it's Steven Levy, called "Cipher"? I can't remember ["Crypto"]. But there are a couple of books that you can read and get pretty up on this. You're not going to be writing your own protocols, but...

Steve: Yeah, and I'm not sure that there's a need for the general public to know more than they know. But the fact that everybody knows that Apple is fighting with the government, you know, that's a good thing for people to know, I mean, to appreciate.

Leo: I mean, that the math is out there, and you can tell anybody, if you're a half decent coder, and I mean half decent, you don't even have to know the XOR XOR XOR trick...

Steve: You don't have to be John McAfee? Is that what you're saying?

Leo: You don't even have to be John McAfee. You could write an algorithm that would give you public key crypto; right? It's not complicated. I think; right?

Steve: You've got to be kind of good.

Leo: You do? Okay, okay.

Steve: Yeah.

Leo: The keygen routines are tough? Okay. Just checking.

Steve: Yeah, it's not something you just kind of want to pull out of your backside.

Leo: Well, the seed is tricky; right? You have to - okay.

Steve: Yeah.

Leo: That's because I'm an eighth decent coder. I don't know. Finally, Bill Barnes writes from Charlotte, North Carolina - and we have to make it finally because in two minute the polls are going to close, and I know you want to rush to the CNN. Bill Barnes of Charlotte, North Carolina wonders about the best drive choice:

Steven, where is the sweet spot in spinning drives between reliability, cost, and capacity? Performance is not a priority in my usage. The home file server - oh, well, that makes sense - is on a single 2TB Western Digital Caviar Green purchased in 2010. Critical files are backed up locally and on Carbonite. Can I safely get a 2-3TB,

or even 4TB, drive from the big box store to keep my data another five years? Would I get better reliability with Windows RAID 1 mirroring and a larger drive or spanning smaller drives?

Blah, blah. Oh, recognizing the value of GRC since SpinRite 2. User of ShieldsUP! since it was in the seven digits. Faithful follower of Security Now! since Episode 15. Programs due to Security Now!: Astaro, Hamachi, Carbonite, TrueCrypt, LastPass, and more. Wishing for CryptoLink and SQRL. Looking forward to you getting back to SpinRite 6.1 and 7. This guy is a hardcore listener. So is there an opt - by the way, we've never talked about it, but somebody released - was it Google? - SSD drive information on wear that was very interesting.

Steve: Yeah, really interesting.

Leo: You need to tell us what it means, though, because I...

Steve: I will. To answer Bill's question, no single drive of any size, and I'm here to tell you, can be counted on.

Leo: You know. This is a man who knows.

Steve: I just built, as I had mentioned before, and my thumb is pointing at it, a monster machine using the Haswell chipset. In fact, I noticed, Leo, that your upcoming virtual reality box, you're using an X99 motherboard and probably the LGA 2011 socket and so forth. I have the same RAM that you guys chose.

Leo: Oh, you're building it, too. This is exciting.

Steve: The Vengeance, whatever it is, although I'm just using it as my workstation.

Leo: Corsair, I think, yeah.

Steve: The Corsair. What I did, this motherboard has one of the M.2 slots, the super high-speed PCIE little tiny card slots. So I got a Samsung 0.5TB, the 512GB SSD. That's my main drive. On the motherboard, could not have a faster interface to the processor. The state of the art, the NVMe interface, screamingly fast. And two 2TB drives mirrored as what the...

Leo: Oh.

Steve: And so here's my logic. First of all, I will never use all of the space on that big NVMe on that big SSD. And I don't want to because I want to take advantage of its wear leveling to move the information around a mostly empty drive. So big things like media

and temporary big downloads and stuff, those go to the spinning media because they're temporary, and they're huge. The idea is this high-speed drive, it'll be hardly 10% full, probably, given my history. I'm running on I think a 30GB partition now in my system, and I'm fine. But the idea is, by leaving it way empty, I'm allowing the wear leveling to have lots of unused space to move things around to, so that it automatically extends its life by the percentage of drive I'm not using. And then, nightly, I will take an image of that and copy it to the mirror RAID.

And I have two 2TB, I think they're 2TB drives, just because I had them around. And they're mirrored because mirroring is faster than RAID, because RAID requires that you have computation in order to do the RAID XORing process. So mirroring is able to be faster than RAID, so I get better performance, and it's a hundred percent redundant. So if one drive dies, I've got the other one. The BIOS supported RAID configuration. I could select AHCI or RAID. I did, so the BIOS understands the RAID. Windows came up; it understood it. And so that's the architecture I chose. For the main workhorse, a deliberately mostly empty, screamingly fast SSD, but nightly images of that drive to a pair of spinning drives that are mirrored.

So, Bill, your answer is no single drive. You've got to have two. I mean, you just do. Because, you know, even the best drive, something catastrophic can happen. The board could fry itself. There could just be something wrong with that drive. So you get - the mirroring is the least efficient because you only get half of the total data storage of the two drives. In a RAID 5, the more drives you add, the more efficient it becomes. And at GRC I'm running RAID 6. I've got four SSDs in RAID 6. So any two of them can fail. None ever has; but I believe in redundancy because, you know, I'm in the drive failure business, and I'm doing just fine.

Leo: There you go. I really like that tip on an SSD. How much capacity should you not use?

Steve: My feeling is you should have storage, bulk storage for, like, media stuff. For example, I've got all of our podcasts are sitting here, available to me. I'm not even sure why. I ought to just spin them off to Blu-ray and archive them. But so the idea would be I will use the SSD sparingly. It's my system drive. It's what boots. I want it to be screamingly fast. And I want to keep it mostly empty. And so apps get installed there; the working files get installed there. But big blobs go to the spinning drive, and images of that every night go to the spinning drives so that, if it ever goes tits-up, I'm able to pull it right back. I've got an image that's never more than a day old.

Leo: You mean bits-up; right?

Steve: Bits-up. That's what I meant.

Leo: Yes, thank you.

Steve: Isn't that what I said?

Leo: Yeah, bits-up. The bits are up. Steve Gibson is at GRC.com. That's his website. Go there to get all the stuff he talks about. Man, this guy is amazing. He's prolific. He's like Alexander Hamilton. He just - scribble, scribble, scribble. But start with SpinRite, the world's best hard drive maintenance and recovery utility. It's his bread and butter. You've got to support the guy. And by the way, you've got to have it. If you've got a hard drive, you need SpinRite. You should also go there if you want the podcast because he keeps a copy, including a unique copy, the 16Kb version, which no one should have to listen to. But if you really want it, it's the smallest audio version of the show.

Steve: If you're counting every downloaded byte.

Leo: There's also a text transcript, which I would venture to guess is the smallest download of all. And Steve pays for those and makes them available, and that's really, really great. I thank you for doing that. We have audio and video on our site, TWiT.tv/sn for Security Now!. And of course you can get it in every podcatcher. This is one of the oldest podcasts in the world. How about that? One of the oldest podcasts in the world, folks.

Steve: Because we're a survivor.

Leo: We're a survivor.

Steve: The other ones have died.

Leo: Yeah, right. Everybody else gave up.

Steve: Digg, that's gone, and everybody else.

Leo: [Wisely], Diggnation, yeah, they just all said, you know, this wasn't that good an idea. But we, we are gluttons for punishment, and we're going to be back next week. And next week I'm not going to be here. I will have been run out the door to get a flight to New York so I can see "Hamilton" next week.

Steve: Yup.

Leo: Spring break for me and Lisa and Michael, our child. And he's a lucky 13 year old to get to see "Hamilton," I've got to tell you.

Steve: Very cool.

Leo: And I'm hopeful that it'll really turn him on the American history and all that. But Father's going to do a great job, and you guys are going to talk about DROWNing.

Steve: DROWN, D-R-O-W-N, yes. A new, really cool hack for breaching TLS communication, no matter, even if you've got 1.2, turns out some guys found a way to get the private key out of that.

Leo: Oh, wow.

Steve: Yeah.

Leo: Oh, it'll all be there. And then, if you want to ask questions, we'll probably do that in two weeks. GRC.com/feedback is the feedback form.

Steve: Yup.

Leo: Steve, thanks so much. We'll see you next week.

Steve: Always a pleasure, Leo, thanks. Have a great trip.

Leo: He always waits for me to say "on Security Now!," but I didn't this time.