Listener Feedback #192

Description: Leo and I discuss the week's major security events 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-464.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-464-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson's here. I'm back. Thanks to Robert Ballecer for filling in the last couple of weeks. We're going to talk about so many interesting topics. Is CryptoLocker neutralized? The Department of Justice says so. And a problem that was a problem with password managers has now been fixed. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 464, recorded July 15th, 2014: Your questions, Steve's answers, #192.

It's time for Security Now!, the show - yes, I'm back. Thank you to Robert Ballecer for filling in for me.

Steve Gibson: Yup.

Leo: Father Robert's a great fill-in. According to some, better than Leo.

Steve: Well, we have - we had a stand-in. There's only one Leo.

Leo: But Tom Merritt when he did the show was great. I think people like hearing a little variety. You and I have been doing the show for almost eight years, so...

Steve: No, no, more.

Leo: More?

Steve: We're, like, we're approaching 10, I think.

Leo: Four hundred sixty-four, it's a lot of years.

Steve: Yeah.

Leo: So people like hearing - it's like when you - if you wear the same shoes for 10 years. Every once in a while you want to put on some loafers.

Steve: Yeah. I tend to, actually, I do wear the same shoes for 10 years.

Leo: Why do I - why am I not surprised?

Steve: Why are you not surprised?

Leo: You have 10 pairs of exactly the same shoe.

Steve: Actually, I do have several pairs that are still yet unworn because I found some fabulous ones. They're not in the refrigerator.

Leo: Share that with us. What is it? What is your perfect shoe?

Steve: I got some shoes - I went to Sport Chalet. And I said, okay, I want the best, like, I don't remember why, but they were, like, running shoes. But I don't run. I bicycle. But...

Leo: Comfy.

Steve: Maybe it was for the stair climber. And I said, I don't care what they cost. And the guy looks at me and says, "Are you sure?" And I said, "I'm absolutely sure because, you know, my feet are important." And I got these things that I called "moon boots" because they were so springy. I mean, they're not air-filled, but they're just foam somehow. And I just kind of bounced around. I felt like I was in low gravity. And for a while I was only wearing them some of the time. But now I still can't remember why I would have been wearing them. But then I thought, why am I not wearing these all the time?

Leo: They're comfy.

Steve: And so that's what I do, yeah.

Leo: So what's the brand?

Steve: They're Nike somethings. They're, like high-end Nike running shoes. And they're just - and so I went back, and I got some more because, you know, Palm went out of business.

Leo: You know, it's a silly thing to talk about, but Dick DeBartolo wears these amazing shoes, these strange Italian shoes. They're also like running shoes. And he just loves them. It's all he wears. And so it just seems like there's something - what it is, it's people who work for themselves. They can wear anything they want.

Steve: That is the case. And I'll never forget when I stood up once and you commented, "Oh, you have no pants on." And I thought, okay, Leo, for our audio audience...

Leo: They're shorts.

Steve: ...it's important for you to explain.

Leo: You have short pants on.

Steve: It's not dangling.

Leo: I don't think I introduced you. Not that you need any introduction. That is Steve Gibson, our security guru at GRC.com, creator of SpinRite. And every week we talk about security. We have some questions to answer. But before we do that, I think we should probably, as we have been doing, not for all nine years, but for many of these years, cover the security news of the week.

Steve: Yeah. Well, actually I think you probably mean the reverse. Because the original concept that you proposed was a weekly newscast about security news. And then it was not long after we folded the idea of feedback and questions in to sort of close the loop. So last week there was so much news, and we just kind of went so long, that we only got four questions in. Although in fairness, the last one was five parts. So it was, you know...

Leo: That's a long question.

Steve: It was a substantial one, from Opher. But I thought, again, we've got a huge

news day, which I wanted to do justice with. There's a big story which has, like, totally dominated my Twitter feed, which was - it's odd because almost 10 months ago, actually 10 months ago, in August, some security researchers at UC Berkeley informed the makers of five different web-based password management tools that they'd found some problems. Four of the five instantly fixed the problems. The fifth one's never been heard from, even though they're just merrily going along with people downloading and using this now known insecure utility.

Then what happened was the news just broke, asking the question whether web-based password managers are secure. And that scared everyone. Well, the fact is, these problems were fixed in September of last year, nine months ago. But out of respect for the researchers, and probably at their request, nobody said anything until now. But they produced a PDF which is 15 pages, two columns, and incredibly interesting and information-packed. That's the topic for next week, is a close look at their analysis of sort of arguably controversial web-based password managers.

The problem is, because it's JavaScript code, can it be secure? And because it's so intimately tied to the browsing experience, which has traditionally had problems with things like cross-site scripting vulnerabilities and all that, what special problems does that mode of operation bring? We know it brings convenience because it's just in the browser, which is where you want your password manager to be, if you need a password manager, and who doesn't because we don't have SQRL yet.

So anyway, so I want to explain briefly that there's no problem, and that we're going to go into great detail next week. And then I wanted to finish up the rest of our questions that we didn't have a chance to get to last week. So we're going to talk about that. And what's really interesting, too, is last week we covered a story about a mesh network of Internet-connected light bulbs and how, because of the homegrown protocol that had been reverse-engineered, it was possible for hackers to get the password for their users' home WiFi networks. Not good.

And so I was bemoaning the lack of standards for the Internet of Things. And lo and behold, we now have three. So we'll talk about that and introduce the new term "sleepy nodes" because you want these things to not burn up your batteries. Many times they're going to be things like smoke alarms and smoke detectors and such that need to run for years on a AA battery. So you need, for that, your nodes to be sleepy. Also, hours ago, Google announced something called Project Zero, which we'll talk about.

The LibreSSL fork of the OpenSSL project came out and stumbled upon leaving the gate. A little coverage of Netflix and Verizon, still arguing, but there is hope. And then a weird report from the Justice Department saying that CryptoLocker has been neutralized. So lots of fun stuff to talk about, yeah.

And I've adopted something that I haven't mentioned before, but I'm trying every week to find an image, a graphic, a something which relates to the podcast and to security and is fun or interesting. And so it's always on the lower half of the front page of the show notes, which you can see it there, which was the really interesting diagram that Verizon published in a blog posting that I link to later in the show notes. What I liked about it was it shows in detail the architecture of Verizon's IP network.

And what happened was one of their customers, you can see him over there in his house in the lower left, happily typing away, he's a FIOS user with a 75Mb connection, of which I am so envious because I'm here in Southern California, and Verizon isn't with FIOS, and boy, would I like to have it. But we're still waiting. Anyway, this guy, apparently, so goes the story, contacted somebody, got somebody's attention at Verizon, and said why is my

Netflix buffering when I have a 75Mb connection?

And so this diagram, it sort of covers, first of all, the internal Verizon network architecture, which is interesting, but also shows what we've been discussing the last several weeks, which is the problem occurs at the boundary, the so-called "peering connection," the peering border between Verizon in this case and the network that they're peering with. It's that...

Leo: We should point out this data comes from Verizon.

Steve: Correct. Correct.

Leo: So take it with a grain of salt. They're at great pains to say it's Netflix's fault, not our fault.

Steve: Well, independent of whose fault it is, I'm talking about the technology here. And so the technology is that it's - and that's what this diagram explains, which is, I think, clearly correct. I mean, no one would argue that it is the border router where the ISPs are peering that you get saturation because of just insufficient bandwidth at that location. So anyway, as we'll see, Netflix has an agreement with Verizon now. And the press is confused because, when Netflix agreed with Comcast, the problem just disappeared instantly. And my guess is that it's just a matter of hardware, that we just need more hardware in this peering location in order to handle the traffic because, as we know...

Leo: You remember the post from Level 3, which is one of these transit providers, that it was the five big U.S. broadband companies that were intentionally congesting the network by not upgrading their end. This implies that it's Netflix not upgrading its end. And that's why I think this is a little deceptive. It's what Verizon wants people to think, or it's what Verizon contends. But Level 3, which frankly also has a dog in the hunt, they would like to say it's not their fault. They say it's Verizon's fault or somebody's.

Steve: Basically, somebody has to spend money...

Leo: Right, right.

Steve: ...in order to create bandwidth for this incredible phenomenon. I mean, if one third of the Internet's traffic in the evening is from one organization, that's a phenomenon. I mean, the whole idea of video on demand still makes my head spin. But it's what people want. Yeah, so, okay.

So last week we talked about the discovery that Google made of fraudulent certificates for their own properties and Yahoo!. And this came about because I'm monitoring the CRLSet now, waiting to see if they're going to block my new revoked.grc.com domain the way they revoked the first certificate I put up. They haven't. They apparently - they're just - they don't care, of they just figure, well, I'll just change it again, as I will. I'm

prepared to do that.

But as a consequence, I noticed, when they updated the CRLSet early last week, then we got, a couple days later, we got an announcement of what that was about, which was that an Indian, as in India-based ISP, I'm sorry, CA, had issued some intermediate or subordinate certificates which were being used to mint fraudulent domain certificates. And Google caught that immediately because, I mean, Google is nothing if not vigilant about the use of their certificates. So what they were doing was, since as we know they're not actually able to revoke things, they pushed out another patch, essentially, to their certificate list, explicitly listing three certificates which were known to have been involved in, essentially were being trusted.

Now, what's interesting is that only Microsoft had this Indian root in its trust store. So Mozilla and Firefox were never at risk. No Apple systems were at risk, either iOS or Mac. And so Chrome was protecting itself from abuse on Windows. And I said at the time, well, Microsoft was going to have to revoke these shortly. And they did, the day after the podcast. So if anyone noted additional updates coming in after the Second Tuesday's updates, which was last week, that's what this was. This was Microsoft adding three certificates to their untrusted certificate list. And anyone who wants to make sure, who's a Windows user, can, if you go into the Microsoft certificate manager, you should see three certificates that are from the NIC certifying authority and NICCA 2011 and NICCA 2014, which Microsoft has now pushed out so that those certificates will not be trusted.

Now, Google then updated their posting with some additional news since then. And they said on July 9th: "The India CCA informed us of the results of their investigation on July 8." So this was the day after last week's podcast. Well, actually it was the day of the podcast, so we didn't cover it on the podcast. And they said: "They reported that [the issuer's] process was compromised and that only four certificates were mis-issued, the first on June 25. The four certificates provided included three for Google domains," only one of which Google was previously aware of, and Google only becomes aware of them when they spot them in the wild, when they catch them being used. So that means there were two others that had been issued, but which Google hadn't experienced. No Chrome browser because Chrome does certificate pinning;, Chrome is able to spot the abuse of Google certificates.

So Google learned of two additional ones, and then also one for Yahoo! that we did know about last week. But then Google said: "However, we are also aware of mis-issued certificates not included in that set of four and can only conclude that the scope of the breach is unknown." So in other words, this Indian CA doesn't have control of their certificates anymore. They were informed of a problem. They said, okay, we'll get back to you. And when they did, they didn't have a full report because Google was aware of additional misuse and may actually have been a little cagey, waiting to verify that this CA, that was naturally embarrassed, as any CA would be that had its trust compromised, didn't have the full story.

So as a consequence, Google has announced that they are going to restrict the domains that they ever trust from this certificate authority, which, you know, this is a mess. This is the last thing you want to do is start having to put special casing code into your browser. But that's what Google's going to do. And this is only necessary on Windows, remember, so not on Android and not on Mac, only when Chrome is being used on Windows because right now Microsoft still trusts certificates from this authority. And Microsoft is now making exceptions and apparently doesn't have all of them because from what we've learned after Microsoft's push, there are additional certificates which are still being trusted by Windows that have not been added to this block list. So I can expect that we'll see additional certs being blocked.

So the certs that - so what Google's going to do is to push out an update to Chrome at some point in the future, and they haven't said when, but I'm sure they will, such that certificates like gov.in will be trusted; nic.in; ac.in, whatever that is. Clearly gov.in we can guess what that is. Bankofindia.co.in will be allowed. And then a few others. So just a handful of domains signed by this Indian certificate authority will be trusted by Chrome on Windows.

And maybe Microsoft will just - it's not clear to me why Microsoft is trusting this now clearly flaky certificate authority where Apple and Mozilla aren't. Which means that nobody using Firefox or any Mac asset or Apple machine could ever be visiting these sites anyway. So it really does sound like this is a place where we ought to just say, sorry, folks, we're not going to trust certificates signed by this authority any longer because they obviously don't have control of their process, and we understand that's something that a certificate authority has to have.

So I pretty much covered what I had next in my notes, which was this issue of web-based password managers. I've got a link to this 15-page PDF, if anyone wants to jump ahead and absorb it. I will be doing that over the next week and completely talk about, as the topic for next week's podcast, what these guys found because this was nine months ago that these problems were found and fixed. LastPass responded three days ago to this news coming out because everyone was worried. Unfortunately, the press miscovered this, as they are wont to do, with headlines like "Critical Vulnerabilities Found in Web-Based Password Managers." Well, yes, nine months ago, and they were responsibly disclosed and by the responsible password managers like RoboForm, like LastPass, and so forth, they were immediately fixed.

And so LastPass wrote, and I think this was Joe and his team, said: "In August 2013, a security researcher at UC Berkeley ... contacted us to responsibly disclose novel vulnerabilities with the LastPass bookmarklets." And then LastPass notes, or Joe notes, he says "actively used by less than 1% of the user base" because primarily that's what you have to do because there are no extensions available currently for Safari, for example, on the iPad. So you either use LastPass's tab browser, or you use these bookmarklets. And they're a pain to use, so obviously a very small percentage of the LastPass user base, less than 1%, uses them. And also a problem with one-time Passwords was found.

These researchers "discovered one issue that could be exploited if a LastPass user utilized the bookmarklet on an attacking site, and another issue if the LastPass user went to an attacking site while logged into LastPass and used their username to potentially create a bogus one-time password." The researchers tested these exploits on dummy accounts at LastPass. And looking at their logs - and remember that LastPass was informed of this immediately, so they were able to look nine months ago, and so they're saying we never saw any evidence of exploitation by anyone beyond the researchers at UC Berkeley.

"The reported issues were addressed immediately, as confirmed by their team, and we let them publish their research before discussing it." So that's why this is coming out nine months later is that this paper was just released. So no one needs to worry. These five password managers, well, sorry, four of the five have been secure for, like, the last nine months. But this is an interesting topic, and so we'll delve into it in depth next week.

Leo: I am curious who the fifth is.

Steve: Yeah. And I'm just not remembering the name. If you look at - if you want to pull up that PDF...

Leo: I'll see if I can put up - yeah, yeah.

Steve: Yeah. If you pull it up, I'll recognize it instantly because it was not a well-known...

Leo: It was not 1Password or RoboForm or LastPass.

Steve: No, no, no.

Leo: I've got the PDF. "The Emperor's New Password Manager: Security Analysis."

Steve: Yes.

Leo: JavaScript always worried me a little bit in these. I mean, that's what makes it open source, but it also means that you can...

Steve: Well, it's not so much open source. It's where the convenience comes from. It's open browser.

Leo: Right.

Steve: And it was - the fifth one was a company called NeedMyPassword.

Leo: Hmm, don't know that one.

Steve: So LastPass, RoboForm, My1login, and PasswordBox are the four that were analyzed. And the NeedMyPassword people never responded. So...

Leo: Did he test others? Were these the only four or five that had problems? Or these were just the five he tested?

Steve: I think these are the five that they looked at closely.

Leo: Okay. So it doesn't mean - you shouldn't assume that this vulnerability doesn't exist with other password managers at this point.

Steve: Correct. And they also mentioned, because I did do some digging, but not enough to give it the kind of coverage I want to, they commented that they're going to be using what they learned to produce their own. So I don't know what the back story is there. But they did immediately disclose...

Leo: That's interesting, huh.

Steve: Yeah, isn't that interesting? Yeah. And they did immediately...

Leo: And we'll talk about this, of course, next week, so...

Steve: Yup, exactly. So the good - what I was bemoaning last week when we covered the story of the light bulbs that were leaking people's private WiFi network passwords was that, with this whole Internet of Things, and now we're seeing the acronym IoT, the Internet of Things, you know, this is something that's very popular. And I mentioned that I was glad Apple announced with the iOS 8 announcement that they were going to be getting into this because I know they'll take security seriously.

And the problem we have with just every random company's wanting to cash in on the popularity of widgets with little microprocessors in them that can participate on people's wireless networks, is that everyone's going to invent their own system, their own network, their own mesh, their own way of connecting these things. And in fact I've not yet ever gotten around to setting up my NEST thermostat, but my buddy has. And he was telling me, like, how he set it up. And he's not super technical, but I couldn't understand from what he told me how it ever knew what his WiFi password was, which worried me.

So the good news is we've gone from me wishing there were some standard to now having too many, in one week. Because of course we don't want many. Or it wouldn't be bad if we had many, so long as they're all secure. And the problem is, I mean, this is like classic, in the early days of crypto, people would just say, oh, look, we're using our own cryptography. It's like, oh, no, no, no, no, no, please don't ask your friend to come up with something that he's really sure is secure and then burn that into your silicon. Use real crypto.

So the problem has been, without standards, and there haven't been any, there's really no choice but to roll your own. So we're seeing, because everyone now really does understand the need, we're going to get standards. So it turns out that a couple months ago there was one announced that I wasn't aware of and that we haven't covered on the podcast. It's known now as the AllSeen, sort of as in all-seeing, but in this case it's AllSeen, s-e-e-n, Alliance. And it's at AllSeenAlliance.org. It uses open source software developed by Qualcomm called AllJoyn, j-o-y-n. And so there is also an AllJoyn.org site. They announced in June that they had acquired their 50th member. And 12 days ago Microsoft joined.

So this is a big group. And their goal, their stated goal is, as they say, "A Common Language for the Internet of Everything." And on their website they said: "Our homes, our cars, and the things around us are getting smarter every day. AllJoyn" - which I guess Qualcomm has trademarked - "is the open source project that lets the compatible smart things around us recognize each other and share resources and information across brands, networks, and operating systems. AllJoyn was initially developed by Qualcomm

Innovation Center, Inc., and is now a collaborative open source project of the AllSeen Alliance. AllJoyn gives manufacturers and developers the tools they need to invent new ways for smart things to work together."

So I've not looked in detail at this. I have a feeling that we'll end up doing podcasts, either collectively or individually, on these three things because this is going to be a hotbed of security interest as these devices connect. And there are two topologies which are possible. One is the traditional star topology, which is what we all have now in our homes with a central WiFi router and our various devices - phones, computers, and things - all connected to it. So there's a one-to-one relationship.

But the intriguing topology which these things will be supporting is a so-called "mesh." And it was a mesh topology that these LED lights we discussed last week are employing, the idea being that they find each other. And a mesh is robust because that's what the Internet is. When we've talked about routing, the Internet is a mesh-routing architecture. It's sort of a best effort. Routers are connected to multiple others, and traffic can find a way. So even if you unscrew a light bulb that's in the mesh, the rest of the light bulbs can route around that one, even if they were using that one for their communications.

So, and of course one of the advantages, to the degree that these devices are really trying to be miserly about power, and that's a huge issue for them because they're typically little smart tags and iBeacons and things running on batteries, they really need to conserve. And we've talked before about how, for example, a smartphone will use up its battery faster when it's far away from a cell tower because it's smart about ramping up its broadcasting power as is necessary for its signal to reach the tower. Similarly, these devices will be able to conserve power if they're closer to a peer and able then to use a lower level of power to get their message across.

The other thing that these are using is technically, or typically, these are very low bandwidth. So they're not having to move tens and 20s and decades of megabits. They're just needing to say, oh, it's 32 degrees at the moment; or, don't worry, there's no smoke being detected here. So those are ultra low-bandwidth, sort of like text messaging bandwidth speed. And that dramatically reduces the amount of power that they need to use for transmitting. The very fact that they're running much more slowly means that they can reduce their power. So all kinds of tricks will be employed in order to keep the batteries of these little gizmos going. And the fact is, the more you have around, the longer they'll individually last because they'll be able to use each other to hop their signal among themselves and eventually over to a router in order to gain remote access.

Number two of these three was announced last week. So these things are just beginning to happen. And in fact the third one was announced today. But number two is the so-called Open Interconnect Consortium, called OIC. And they've got, again, a lineup of major companies. They said: "The Open Interconnect Consortium is being founded by leading technology companies with the goal of defining the connectivity requirements and ensuring interoperability of the billions of devices that will make up the emerging Internet of Things." And among the members are Intel, Broadcom, Samsung, Atmel, Dell, and Wind River Systems, with more expected soon. And this is a young group, but people are joining quickly. So that's OIC, OpenInterconnect.org.

And then finally, in this morning's announcement, is the Thread Group. And that's just ThreadGroup.org. And these guys are a little different. First of all, Thread is a protocol based on IPv6. And it's what Google's NEST group is already using, that is, NEST Labs at Google is using this protocol. Samsung and ARM and Freescale, which is the embedded component that used to be Motorola, something called Big Ass Fans - and I was curious,

so I went, and that's what they sell is, in fact, their home page talks about how our fans are smart. And it's like, okay, well, so they're large and smart fans.

And then Silicon Labs and Yale Security are among the current members of this Thread Group. What they've got is a so-called Thread Protocol. And so they distinguish themselves from the other two because they're not an IOT platform, but a wireless protocol which the other two platforms, AllSeen and OIC, could work on top of. So these guys are, again, this is where they talk about Thread is their protocol, which I'm sure we'll be talking about in the future, which supports so-called "sleepy nodes," which are able to operate for years on a single AA battery. And there is an IEEE standard, it's 802.15.4, which is specifically for the so-called 6LoWPAN, Low power Personal Area Network protocol. And it is all IPv6 based.

So again, I've not had a chance to delve in depth into these. We'll sort of keep our eye on them and see which ones end up succeeding. But the good news is they really are focused naturally on security. And in fact, I don't remember which one it was I ran across - and I'm not seeing it in my notes. Somebody specifically said that they were focused on - oh, yeah, it was the OIC group said they would be focusing on security and authentication, initially targeting the home and office, and then automotive applications. And of course the Thread Group with Google's NEST Labs and the others, they're using Thread Protocol now. And I'm glad to see that it's open, and everybody is looking closely at security. And now that Thread is open, I'll be able to answer my question about how the NEST thermostat got configured to the network and found out what was going on.

Google also this morning - it was a big morning for news - just announced Project Zero. And I like the way - it was Andy Greenberg writing for Wired, and I'm paraphrasing what he wrote. But, for example, let's take the case of young American hacker George Hotz, who we've discussed a few times on the podcast. Back in '07, at age 17, George famously hacked and cracked AT&T's lock on the iPhone. Later he reverse-engineered and cracked the PlayStation 3, which caused Sony to sue him. And they settled with George agreeing never to hack another Sony product. Then he cracked Chrome's OS security, winning himself $150,000 Chromium Award from Google for doing so. And then, two months after that, Chris Evans, who's - I can't remember what his business card says. I saw him referred to as the "Hacker Herder," but that's not what his business - I think his business card says "Troublemaker" as his official title. Anyway, they hired George Hotz full-time to join Google's new team of elite hackers in Project Zero.

So this has been characterized various ways by the press that immediately picked up on this press release and blog posting. But, for example, Tavis Ormandy, whom we've discussed often because he's been a great source of revelations of security problems, he's been recruited by Google. He was already at Google, but now he's part of this team Project Zero. They will shortly have or soon have 10 members on the team. And they are full-time employees of Google, all of them. And Google has said, "We're hiring." So if you are a talented hacker who enjoys Internet security puzzles or have had these discoveries to your name in the past, Google would like to hire you.

And the focus of Project Zero is wide open. It is meant to be for the good of the Internet, to find vulnerabilities which will be responsibly disclosed in anything. Doesn't have to be Google properties. It could be anything. And Google has said they will responsibly disclose them, then publish them, once patched, to an external database so that everyone can see what Project Zero is doing. So this is just another initiative inside Google to push things forward and improve the security of the Internet. So I think it's great.

So we talked about the forkings of OpenSSL. We've discussed of course famously the big

Heartbleed problem with OpenSSL, and that LibreSSL was the first fork of that. Well, it was released, what they call v2.0.0 was released last Friday. And it was shortly tweaked for some portability issues which were immediately discovered. In fact, 2.0.1 was then released on Sunday, so two days later. So that means that immediately after the release, people took the source and began compiling it in their own environments for their own use and immediately found some problems that the Libre guys fixed and released 2.0.1.

So here's the problem. A blogger, Andrew Ayer, wrote: "Despite the 2.0.x version numbers, these are only preview releases and shouldn't be used in production yet, but have been released to solicit testing and feedback. After testing and examining the codebase, my feedback," writes Andrew, "is that the LibreSSL PRNG," which we know is the Pseudorandom Number Generator, which is crucial, "is not robust on Linux and is less safe than the OpenSSL PRNG that it replaced."

Now, the problem is interesting. The way Unix servers have traditionally handled multiple connections is through a process known as "forking." You'll have sort of a root server which accepts connections. But the server code itself is not multithreaded by design and not able to handle multiple connections in a single instance of code. So when a new connection is presented, it forks itself, which is the original Unix term, meaning that it just makes a copy. It, like, clones itself. And so that fork, that clone of its code is given that connection to handle. And when another connection comes in, another fork is made.

And so the idea is that, rather than a single base of code being able to simultaneously, typically through multiple threading, handle multiple connections, individual copies are made per connection. That's really not the modern way to do it. And the newer server architectures on Linux and Unix platforms are not forking; they're using more advanced I/O protocols so that - because there's lots of various sorts of inefficiencies. Well, imagine if you have a pseudorandom number generator which is deterministic software based, so that it's got an entropy pool. And that entropy pool is driving an algorithm which is producing random numbers. And you make a copy of it. Well, both copies are then going to generate the same random numbers. And that's what's happening.

Now, the designers of this PRNG technology, first on OpenSSL, understood that this was a problem. So what they do is they try to detect forking because there's something in Unix known as the PID, the Process ID. And there's a guarantee that the child process of the parent, which forks to create the child, will have a different process ID. But it turns out that that guarantee does not extend to a grandchild, that is, the guarantee is only a child-parent relationship. And so that means that, if you have a second child-parent relationship, like a second-generation relationship, the grandchild of the original process can have the same process ID, which would mean that it might not detect that it had been forked, and so you'd end up generating the same pseudorandom numbers until there was a reseeding.

Now, users of OpenSSL who are on the ball understand this. So they deliberately reseed after a fork in order to not have the problem that a clone has been made of a pseudorandom number generator which is just using the data in the clone to generate future random numbers. First of all, this is a horrible thing to be doing. We did a podcast some weeks back where I talked about the architecture that I designed for the entropy harvester in SQRL that has none of these problems. And there's no reason any pseudorandom number generator today could not be well designed. But not only is OpenSSL not well designed, but Libre cloned it. So what happened - but also broke it in the process.

So under OpenSSL, if you understand the danger of forking in order to handle additional connections, the first thing you do after forking is you reseed your pseudorandom

number generator. And there is a call to specifically request a reseeding now. Before control returns to you from that call, the pseudorandom number generator will be reseeded. Turns out the LibreSSL people, for unknown reasons, NoOp'd that call. So they have neutered the programmer's ability to recognize, after a fork, that there are synchronized separate pseudorandom number generators, both generating the same numbers, until they reseed for whatever reason. So it looks, based on this, it looks like, as Andrew said, this is not ready for prime-time yet. And I hope that the LibreSSL folks will give this some time. And really OpenSSL ought to seriously look at improving the design of their pseudorandom number generator. As we know, crypto depends on unpredictable pseudorandom numbers. And this is clearly crazy that this is the way the system is operating.

And I had next in my notes Netflix and Verizon, which we also already discussed. Verizon did say that, in their posting, and I'm trying to skip the stuff we - they said: "Therefore, we are working aggressively with Netflix to establish new, direct connections from Netflix to Verizon's network. This doesn't prioritize Netflix traffic, but it ensures that their traffic gets on our network through direct connections, not middleman networks, that are up to the task. The benefit of these direct connections will be twofold. First, Verizon customers who use Netflix will have a significantly improved experience as Netflix traffic flows over non-congested links. Early tests indicate that this is the case. The other benefit will be that the congestion that we are seeing today on those links between these middleman networks and our L.A. border router will likely go away once the huge volume of Netflix traffic is routed more efficiently. This will improve performance for any other traffic that is currently being affected over those connections."

So that's what we were talking about when we explained that, if a border router is saturated, and there's no traffic prioritization, not only is the bulk traffic which is creating the congestion unable to get through, but with the way routers handle their buffers being overfull is they're often, you know, they're just discarding whatever data can't fit in the network interface controller's, the NIC's buffer, and so other traffic is having a problem.

And so the point is, what Netflix is saying is that - or, I'm sorry, what Verizon is saying is their agreement with Netflix will be resulting in a direct connection between Netflix's network and theirs, and not going through middleman carriers. And the point is that then that will free up the border routers in the existing peering agreements, and everything should work right. So that's why I said looks like there's hope here. It's just probably taking a while to get the hardware and everything in place and configured, and maybe some T's dotted and I's crossed first.

Leo: That's basically what a peering relationship is, right, is that you put your - we have a direct connect.

Steve: Yeah. What I don't understand, and I've looked around, and I still don't quite understand, but they talk about a relatively balanced flow of traffic. That term keeps coming up over and over.

Leo: And of course that's nonsense because there's nothing Verizon customers would send to Netflix that's anywhere near the same amount of traffic that Netflix sends to Verizon. It's an asymmetric relationship.

Steve: But that's new, Leo. That's the point. The traditional peering relationships have

always relied on a balanced flow. And in fact I remember when I set up my deal with Level 3, they wanted to know what my inbound and outbound bandwidths were. And they're not happy when they're really out of balance. And I don't really, I still don't get why that matters. I mean, but it does, because even earlier in that Netflix blog posting they talked about a balanced relationship. As if...

Leo: Well, the way you balance it, if you can't balance it with bits, is with bucks.

Steve: Correct. And the idea was that, traditionally, peering relationships would have a near symmetric flow. And somehow the idea was that the people on both sides would thereby be getting equivalent value from peering, which - and I still don't understand why that's the case.

Leo: It makes sense if it's two Internet service providers connecting. It doesn't make sense if it's an Internet service provider and a content provider, like you, me, or Netflix.

Steve: Exactly.

Leo: It's nonsense. Because...

Steve: Although, you know, Netflix is getting value because they're offering content for which then they're being paid for offering content. Verizon is getting value because, exactly as Brett Glass said, the number one question his numbers have is can I get Netflix? So for Brett to be able to say yes, for Verizon to say yes, that's providing - Verizon is an ISP and delivering bandwidth, they're not delivering the content, but they're delivering the bandwidth that allows the content to get to their customers. So everybody's getting value from - even though it's an asymmetric bandwidth flow, the fact is that's today's world. As you say, we now have content providers. And you and I are doing that right now.

Leo: Right.

Steve: We're providing content from a network out to a bunch of listeners.

So the Justice Department released what was sort of a press release that got picked up by a number of magazines. And I don't get it. This was on Friday the 11th. They said that the Justice Department has reported that CryptoLocker has been neutralized by the disruption of its network and cannot communicate with the infrastructure used to control the malicious software. As a result, CryptoLocker is effectively nonfunctional and unable to encrypt newly infected computers.

Now, part of that I understand because we do know that, even if you're infected with CryptoLocker, no encryption can take place until CryptoLocker is able to access a key server and obtain the public key which it then uses for generating a key such that that server holds the private key, and you have to pay them in order for it to decrypt the symmetric key, which is used for encrypting all your files. So if the infrastructure is shut

down, and if that's true, then CryptoLocker, even new infections, cannot take action, essentially. They're reaching out and trying to get to their infrastructure, which the Justice Department is claiming has been shut down.

So I have not been able to find any further evidence of this. SC Magazine, the IT security magazine, reported that in its nine months of existence, the CryptoLocker ransomware extorted more than $27 million from victims. Of course we've often covered CryptoLocker and the variations of it, and Bitcoin and, you know, was it Western Union? Or, no, it was something that you could buy in a store, and they were, like, out of them because the CryptoLocker infections were so prevalent. And this is BitDefender's analysis, having watched this. And they're the guys that were the original discoverers and who mapped the CryptoLocker network. And apparently CryptoLocker saw more than 12,000 victims in less than a week as it was initially spread through phishing email.

So even the BitDefender guys said that, while it may be nice that CryptoLocker, like the existing infection has been disrupted and cannot get a hold of its infrastructure, anything that makes $27 million for bad guys in nine months is going to be, well, we've already seen some clones of it. But they're expecting variations on CryptoLocker that will avoid whatever has been done in order to disrupt their network. So I think we need to consider this is maybe a brief respite from CryptoLocker, and we have not seen the end of it.

And I just wanted to mention I have two miscellaneous topics. I was watching you yesterday with iPad Today. And that Osmo, O-s-m-o, I thought was so clever. That was an iPad app where you put an iPad on a stand, and you put sort of a little hood on the top of the iPad...

Leo: Wasn't that cool? Yeah.

Steve: Yeah. And what that is, of course, it's a mirror that aims the camera down at the surface in front of the iPad.

Leo: So the camera's...

Steve: Yeah, exactly. And so I just thought that was so clever. So the idea was, just for the people who didn't see iPad today, it's called Osmo. It's meant for, like, kids. Yep, and now Leo's showing it.

Leo: I'm showing the video, for video viewers, yeah.

Steve: Yeah. And so the idea is it connects your iPad to the real world. So you can, like, play with Tangram things, like move the tiles...

Leo: They show a Tangram on the iPad screen. And then the kid has to do the same Tangram. And the camera sees it and knows when it's right.

Steve: Exactly.

Leo: And it can say, oh, you got the bunny rabbit, good job.

Steve: Yup. And in fact you're even able to, like, draw pictures of physical systems. And then it will pick up the drawing on the screen and animate the physics of it.

Leo: Isn't that cool? It's so cool.

Steve: Yeah. I thought it was such a clever, again, I just loved the elegance of that, that just by putting the iPad on a stand and doing something as simple as a little mirror so that the front-facing or the user-facing camera now is looking at the desktop. Just-just really clever. And you can do things like with Scrabble tiles and all kinds of stuff.

Leo: I like the physics for the drawing engine is really neat. I guess they have three different devices or games that they're offering. So it's really cool.

Steve: Not inexpensive. I was a little shocked by the price.

Leo: Yeah, especially since now that I know it's just a mirror...

Steve: Exactly.

Leo: It's really you're paying for the app.

Steve: Yes.

Leo: Yeah, anyway.

Steve: Also, number two topic, I just wanted to mention for our listeners that "The Strain," which premiered on Sunday on FX, was very fun. IMDB gives it an 8.6. And I got a kick out of the L.A. Times review yesterday. They said: "'The Strain': These vampires aren't sexy; they're just deadly." And they said: "The vampires in FX's new thriller 'The Strain' are not - repeat, not - romantic. They're not brooding or conflicted or passionate or sparkly. They do not pout, pose, or toss off come-hither glances. And not a single one of them looks anything like Alexander Skarsgard." And he, of course...

Leo: "True Blood," beautiful vampire.

Steve: Yes, yes. He plays Eric Northman on "True Blood." And anyway, so this is from Guillermo del Toro. And they said: "As fans might expect..."

Leo: Oh, he's great.

Steve: Yes.

Leo: Oh. If he's doing it, then I'm more interested, yeah.

Steve: Yes. It is his. It is his. They've been teasing it all, like early summer as starting in July. And I've just been, I mean, like really neat teases. And it's like, come on, come on, you know. So we're finally here. And so the L.A. Times says: "As fans might expect from creature-creator extraordinaire Guillermo del Toro, the undead in this horror series are truly terrifying. They're also parasitic and viral, the product of a kind of contagion that single-mindedly seeks out hosts."

So for anyone who missed it, it is reairing, the premiere episode, which is titled "Night Zero," is reairing this Thursday, Friday, and Sunday on FX, prior to the airing of the second episode, which is titled "The Box," which airs next Sunday. So, and I enjoyed it. I mean, it looks like we're going to have a fun series. So that's neat. And as you said, I mean, it just looks like it's well done.

I'm continuing to work on SQRL, of course. In fact, this is - I'm holding up to the camera a rescue, a so-called "rescue code" that was generated and printed by SQRL, by the Windows client of SQRL a couple days ago. So that work is proceeding. I expect to have all of the identity management stuff finished within a couple days. And that's completely separate from the protocol. So I will turn the client loose to the denizens of GRC's newsgroup, the SQRL newsgroup, for them to pound on and play with and find any mistakes that I made while I'm working on the protocol side. And somewhat excitingly, when that's done, so is SQRL. So, you know, making great progress.

Leo: That's cool.

Steve: Yeah, yeah, yeah. I'm very excited. And while you were in Hawaii a German student, Ralf, and I can't remember or pronounce his last name, he's doing - I think, no, it's before you left because I remember telling you that he had done his master's on SQRL. And he showed his boss SQRL running on his Android phone, and the boss was just stunned and was going to go out the next day or after the weekend and buy an Android device just so he could play with SQRL and show it to their clients because they're going to support it. He's given Ralf a budget, and they're going to be bringing up an iOS version of the client also.

Leo: Awesome. Awesome.

Steve: So, beginning to happen. And I did see in my mailbag a note, actually a question that I wanted to explain, from Dan Hankins in Scottsdale, Arizona, who was wondering about SpinRite operating in a virtual machine. He said: "Much thanks for Security Now! and SpinRite. Love the show and the product. I have a SpinRite question: I recently started running SpinRite for Level 4 maintenance from within a VMware virtual machine on my Linux host. I discovered, much to my surprise, that running that way was more

than an order of magnitude faster than native. Why would this happen?

"I am concerned that, because I have write caching turned on, the maintenance pattern writes are not actually ever being written to the hard drive. If that's true, that would defeat Level 4's attempt to refresh the surface. That would make Level 4 equivalent to Level 3 or 2. Should I turn write caching off, or is something else going on?" And the answer is, yes, something else is going on.

What's happening is that many BIOSes are not supporting Ultra DMA mode. They just support what's called "Programmed I/O," PIO. And SpinRite 6 uses the BIOS for performing its bulk data reading and writing. It drops into talking to the drive directly for recovery because the BIOS doesn't give us nearly enough control. But it does use the BIOS for bulk I/O. But SpinRite does also disable write caching itself so that it's actually doing reading and writing, which is why the BIOS, which is stuck in PIO mode, is so slow. So what's happened is in VMware or in VirtualBox, and in probably all the VMs, have much more state-of-the-art BIOSes. And they're, like, a virtual BIOS. But they're supporting Ultra DMA.

Now, so what that means is that, even though SpinRite, running in VMware or VirtualBox on Linux or Mac or PC, even though SpinRite is still using the BIOS interface under SpinRite 6, the BIOS is then using Ultra DMA so that that's not slow. Now, it's not as fast as 6.1 will be because it's still using SpinRite's small track size buffers, which is what it's traditionally used. And when I've been talking about SpinRite 6.1, one of the things, the technology already finished for 6.1 is that I'm using maximum transfer size buffers, 32MB buffers. And it just, really, it screams. So anyone using SpinRite 6 today who is seeing it run really slow on their machine, if running in a VMware virtual machine or in a VirtualBox, which is free, virtual machine is an option, it may very well run vastly faster for you. And you're still getting a full benefit of SpinRite.

Leo: Hey, I just got this in the mail. I thought you might be interested in seeing it. My very first chip-and-pin credit card.

Steve: Yeah, heard it's coming.

Leo: This is a MasterCard. Yeah, supposedly by next year. And I had to login to activate it. I had to give it a four-digit PIN number. And of course our international viewers, this isn't anything special. They've seen these for years. But this is new to the U.S., the idea of putting in a chip. I don't know what - do you know what the chip does in the chip-and-pin? Is it like the VeriSign dongle? Is it like - or is it just NFC?

Steve: I have not looked at it. Yeah, I have not looked at it closely yet. So what we're going to have, although it is a contact reader, so we're going to have to have readers changing from mag strip to essentially contact readers, so you'll stick that into a slot and then enter your PIN. Certainly I need to spend some time and come up to speed, and we'll explain it to everybody because...

Leo: It might simply be a memory chip with a number in it or something, a token, yeah.

Steve: Yeah, but certainly better than a - I think they did something better than it just being a ROM because we already had that on the mag stripe.

Leo: You're right. It's a good point.

Steve: So I think it participates more actively.

Leo: It'd have to; right? What would be the point otherwise?

Steve: Yeah. Cool.

Leo: Yeah. It's finally happening. And these will make you much more secure until...

Steve: And yours also has a mag stripe on the back?

Leo: Yes.

Steve: Yeah, so...

Leo: Somebody's saying there may also be a touch-to-pay RFID in here. I don't know if that's the case. It doesn't mention that, but who knows.

Steve: Yeah. It's neat.

Leo: Ladies and gentlemen, we have some questions for Steve because we didn't get to them all last week. Starting off with Christopher Hunt, who wonders why not use "pass=openwireless.org" for the OpenWiFi SSID: In listening about and further researching the Electronic Frontier Foundation's openwireless.org site - that was the mesh network they were proposing, or are proposing - they suggest those who support or join to rename their AP SSID to "openwireless.org." Would it not be better for all involved to use a broadcasted, non-suppressed SSID of pass=openwireless.org, with of course a password of openwireless.org? That way the AP is still open, yet it isolates all users involved. In other words, he's saying turn on WPA, but just give out the password.

Steve: Yeah. I think once we see this in operation it'll make more sense. The EFF is in a conference, like in the next couple days. And I keep - the conference is not one I'm familiar with, so I can't - I keep forgetting the name of it, although I've seen it several times. And in fact the movie, remember the hacker movie dot - I don't remember if it was dot com or dot org [TheHackerMovie.com]. But it was one that we talked about that was going to be made by a private filmmaker. He ended up releasing that and making it free for 24 hours recently, and he'll be airing it at this conference, which is where I saw the name again recently, although I can't remember it.

Anyway, the point is that EFF is making firmware available for a specific router which essentially gives it this capability. They're promoting the idea, as we talked about on the podcast a few weeks ago, which is what stimulated Christopher's question, of a sort of a two-facing router. You'd still have your own private password running local encrypted communications. But you would also have another, sort of like a guest mode. In this case, though, it would - and I don't know what their SSID is. I think the SSID is wireless, in fact now I do remember that it's wireless.org is the SSID you're supposed to use.

And my point is that this works with a certificate in your device. So devices can have certificates. And so there would be a certificate for this network that would allow it to encrypt your communications without needing to manually input any sort of a password. And that certificate in your client device would match the certificate that would be burned into the firmware which EFF is providing. And the point is it gives you encrypted wireless with zero effort.

And I ran across an odd story in the last week, talking about - I can't remember how I encountered it. But it was a posting about how patrons of a restaurant were complaining about the service in - I think it was in, or, no, I guess it was now, in 2014. And it happened that the restaurant had cameras just for general dealing with problems that occur in the restaurant purposes, from '04. They still had the recordings from '04. So they were able to analyze what people do in restaurants now versus then. And the story is that people are so hooked on their smartphones that their food is getting cold, and they're, to a much greater degree than they were before smartphones, they're complaining and sending it back. The waiters are coming by, but they haven't looked at their menus yet...

Leo: Oh, lord.

Steve: ...because they're busy with the smartphones.

Leo: On their damn phones.

Steve: Exactly. And they're having trouble, they're complaining that they can't get on the restaurant's WiFi and asking for instructions for how to get on the WiFi. Meanwhile, of course, no other tables are being served, and people's food is getting cold, and the waiters are having to come back because of phones, smartphones.

Leo: WiFi. It's what's for dinner.

Steve: And so the idea would be, if establishments adopted this, and if this became popular, people would have downloaded the matching certificate, which would be freely available, and their phones would just simply be on the Internet. WiFi would just become transparent. No need...

Leo: Better yet, get a cell phone jammer and let them eat the GD food. All right. I'm sorry.

Steve: Yeah.

Leo: Yeah. I guess there's no going back.

Steve: Well, jammers are legal in Las Vegas, and...

Leo: Oh, are they, really?

Steve: Yeah. And famous...

Leo: Oh, I guess in the casinos, sure.

Steve: Exactly, because you don't want people getting up to no good in casinos. And so casinos block cell phones. But it's very illegal to do that...

Leo: Sure, because the babysitter, you know, you might be expecting that you could get a call from the babysitter and not know...

Steve: And 911 calls get blocked and so forth, yeah.

Leo: We're just going to have to live with the fact that nobody's paying attention to anything anymore.

Steve: So we will keep our eye on this EFF initiative and see how it goes, yeah.

Leo: It's kind of the contrapuntal response to the Comcast-Time Warner thing. Turn your access point into a public access point.

Steve: And they mention that this increases your privacy because it gives you plausible deniability. You're saying, hey...

Leo: I don't know.

Steve: I'm making my wireless available. So that wasn't I who did that. Could have been anybody.

Leo: Yeah. Question 2 comes from a chatter, Jim P. in Chicago. He wants to know whether LastPass's "Warn before filling insecure forms" actually works. Before I left we were talking about this problem.

Steve: Yeah.

Leo: In this week's show - and I think it was a couple weeks ago - you mentioned LastPass has a feature that will warn you when filling insecure forms. But I've tried it on two sites that have insecure registration forms, WinSuperSite.com - Paul Thurrott's site - and ChannelPartnersOnline.com. And I didn't get a warning from LastPass even when I submitted the insecure form. Are you sure it works?

Steve: Okay. So that's a great question. And that was a popular tip, so I wanted to address this. First of all, anyone who's curious can go to TiVo.com because I'm a TiVo user, and every time I go to TiVo.com that warning pops up to warn me that TiVo's form submission is insecure. And it also turns out that the NYTimes.com blogs page, so NYTimes.com/blogs, same thing happens. Jenny is a big New York Times reader, and so she's often sending me links. And I was following a link to some story that she wanted to share with me, and I couldn't get there, so I dug around and went to the blogs page, and up popped another warning. So I know that it can work.

But Jim's point is, is it guaranteed to work, I think. And the answer is no because of scripting. When a form is there, LastPass is clearly looking at the page content. And we know that it's doing that because it's a form fill-in application. So it's looking, it's parsing the HTML of the page and understanding what's there. And it's looking to see whether the action URL for the form is HTTPS. Unfortunately, in super fancy sites like Paul Thurrott's WinSuperSite.com, it's not simple HTML.

And in fact I was curious about Jim's claim about Paul Thurrott's site. And so I went there. And oh, my lord, I have no idea what is going on. I mean, I captured the traffic. I clicked on a form. And I don't know where Paul got the gobbledy-gook that he's serving, but it is amazing. I don't know what it's doing. But somehow he's pulled things in from different places and...

Leo: Don't blame Paul. He hates their content management system. It's the company, Penton Media, that runs it.

Steve: Well, good.

Leo: He's not a fan of it, yeah.

Steve: Oh, boy. So the point is that, if scripting is involved - and, boy, is it involved because, if you don't have scripting on, then nothing works there. So there just isn't obvious form content. Something is somehow being rendered on the fly, being pulled in. In fact, I think I got like a pop-up-like thing floating over the top of the page, and only after I completely lowered my defenses and told NoScript to trust 26 other things. And it's like, oh, boy.

So the answer is I know that it can work. But due to the nature of scripting and how crazy things have gotten, you certainly can't - it can't be guaranteed to work. But it does pop up when I run across sites from time to time. And I'm always happy to see it because, if it's credentials or credit card information or something sensitive, you definitely don't want to be submitting that on an insecure form. In fact, there was a site,

I can't remember now what it was, just the other day where I was really uncomfortable. I mean, I was glad that I had that warning. So for anyone who didn't take action, by all means do if you're a user of LastPass because, although it may not pop up in every instance, when it does, it's definitely good to know.

Leo: It's not in the settings, it's in the preferences. So you'll open LastPass, and you'll change your preferences in the advanced section, "Warn before filling in insecure forms." And don't forget to save it. I forgot to save it. So it's not the default.

Steve: Yeah. And then go to TiVo.com, Leo.

Leo: Oh, let's try it, shall we?

Steve: Yeah.

Leo: Shall we try it? I have a TiVo account, so I can log in on this insecure site. Let's just - so first I have to say - oh, I'm signed in, so let's sign out. Sign in. It's going to take me to the sign-in, well, see, it remembers my - that's too bad. I'll have to delete the cookie for that to work. Or I could do it in Safari.

Steve: Anyway, so...

Leo: Go ahead.

Steve: For what it's worth...

Leo: It's worth doing, even if it's not a hundred percent.

Steve: ...it pops up. Yes, it's definitely worth turning it on. It has saved me several times when I've been at sort of off-the-beaten-path sites that are saying, oh, log in. It's like, oh, really going to...

Leo: Man, it's all right. Everything's okay.

Steve: Yeah. So what LastPass is doing is it's looking at the page. And it is an annoying aspect of HTML that we can't see what it is that the form submission is doing. The page we're on can be insecure, yet the form submission can be secure, and vice versa. And in fact I think I was on a secure page, and LastPass was saying this form submission is not secure. It's like, huh? Why would they do that? But LastPass saw that it was a nonsecured submission.

Leo: Something was going on. You know, it didn't warn me. But never mind.

Steve: Yeah. For what it's worth, I see it all the time.

Leo: Good. Yeah, and there's no reason not to turn it on, even if it doesn't work a hundred percent of the time.

Steve: Yes, exactly.

Leo: Question #3, Markus Mix in Lindlar, Germany, which he says is near Cologne, or Kln, wonders why the world thinks LastPass is rock solid. Hey, why does the world think that? I'm a listener of Security Now! for three years. Most often I agree with your much appreciated opinion, but sometimes I cannot follow your arguments. I'm writing because I can't follow the argument about why LastPass is secure. It can't be proven because it's not open source. In Episode 256 - some moons ago - Steve said he tested for some days and has come to the conclusion LastPass is secure. Well, that's not the normal TNO gold standard. Am I wrong? Am I missing something?

I've been using Keepass for years - which is open source - but now I'm looking for a secure password manager which supports per-person team-access controls to get away from different Keepass databases and move to a centralized solution like LastPass. But it needs to be TNO gold standard, especially when using a cloud solution. I don't like it because I don't trust the public cloud. I would much appreciate any help about what I missed because I really want to understand why you trust LastPass.

Maybe after the NSA revelations it would be worth checking back on the subject of LastPass in a Q&A episode in the future? We all love your podcast and hope that we can hear it for a very long time. Best regards from the world champion Germany Lindlar - goal. Okay. I added that. He didn't say that.

Steve: So we would like security to be black and white. And if there's any lesson we've learned, it's that it's not black and white. Keepass is open source. Are you aware, Leo, is there an open source cloud-based solution that Markus would like?

Leo: Keepass is the only open source solution I know of.

Steve: Yeah. So...

Leo: Now, but you've got - you looked at the JavaScript. So you were able to see some stuff; right?

Steve: Yes. Well, what I did was, the reason I use it and trust it, and it is the most popular password manager in the world today, is that they have been very forthcoming about the technology and the protocol. That is, Joe laid out exactly the way it works. And

I and other independent coders coded up an implementation for ourselves that generated exactly the same data that LastPass is sending to them. So I know, I mean, it's like it's better than open source. It's independently written and verified identically operating. Which is, like, beyond auditing the source code.

And that's why I'm excited about many people implementing SQRL is that, when all of our stuff matches, we're all, like, cross-verifying each other's work. There's just nothing better that you could do than independently write to a spec and have the same thing come out. Then you're verifying everybody's interpretation. So...

Leo: It doesn't mean there's not a backdoor, though. Because you could have a backdoor and still produce the same result.

Steve: Absolutely. And, well, let's see. Okay. So the fundamental problem with - and this is where we get into tradeoff - with a web-based solution is that we are, when you use LastPass, and you have the plugin, you're getting the code, your browser's getting the code from them. That allows them, as with Google, for example, to be updating themselves. That's why nine months ago, when these researchers at UC Berkeley found a problem, Joe was able to fix it in a day, and we all got the benefit of it the next day.

So, yes, it's true, the flipside is that the software is dynamic. That is, we're always receiving it. But that's true with Windows and Microsoft's updates. It's true with Google and Chrome. I mean, that's the way the world is today. So the only way to know for sure is write your own. Or find something open source and lead it or implement it so that it's compatible, and then never change. I mean, there are, if you want absolute security, you have to start somewhere. And you have to decide where you're going to start trusting. There have been people who've wondered if Intel architecture, if the Intel chips don't have a backdoor in them. Well, go get some sand and melt it into silicon, make it very pure, and make yourself a processor. Or I guess you could just take an FPGA and start there. You really don't have to start at the silicon.

But my point is, the fact is you are already trusting a huge infrastructure. You're trusting certificate authorities. We don't know how many of those actually are NSA fronts. I mean, we're already trusting. So you've just got to decide where you feel comfortable and where you don't. I'm very comfortable trusting LastPass. I think, you know, could I be wrong? Yes. Am I probably wrong? I see no evidence to believe that.

Leo: And as we've learned from TrueCrypt, even if the source code is readable doesn't mean that anybody's validated it. And it's quite an...

Steve: Or OpenSSL.

Leo: Or OpenSSL.

Steve: OpenSSL and Heartbleed. It was in use for years and with a big problem. So, yeah, the very fact that it's open is not - really what you want to do is you want to have independent developers write code to a spec and have the results match. And that's what many people did for the protocol that LastPass uses to encrypt our stuff in the cloud. We understand that we're taking our email address and our password, how that's being

hashed to create the key for encryption, and then that's being hashed, and that creates a tag that LastPass uses to identify us. Yet because the tag is on the other side of a hash, and hashes are by design not reversible, they can't go backwards and figure out what it was that we hashed to get the tag. We did cover this on that Episode 256. That was probably the LastPass podcast that I did, as you said, like five years ago, where I said, okay, this is what I'm using. I understand this.

Leo: And based on that, I've been recommending it and using it myself ever since. In fact, we use it for our enterprise password management, and we've offered it free to every employee to encourage them to use it themselves, offered a personal version, and I'd pay for it.

Steve: Jenny's using it and says, oh, my god, now I can finally get back into the sites I got into before because...

Leo: I was surprised, you know, we sent out a memo saying anybody who wants LastPass for your personal use, we will pay for it, and nobody took me up on it. Which either means they all are using it, that's what I'm hoping, or they don't care. But I think our fine staff probably all knows, they've heard me rail on and on long enough, they probably all were using it.

Steve: Well, and the free version really does as much as you need, too.

Leo: Well, the only thing you get with the paid version is, I think, mobile. And since most of us now use our phones, having LastPass on the phone...

Steve: Oh, my god.

Leo: ...is really a big part of it. It's only 12 bucks a year, it's not...

Steve: Well, and we also like the fact that it has a clear economic model. I'm happy.

Leo: Yeah, I don't - yeah.

Steve: I'm happy to pay Joe a buck a month in order to know, you know, I get no ads and no funny business.

Leo: And I should point out that life is full of trusting. I mean, if you drive down the street, you're trusting that the guy going the other direction isn't going to swerve into you head-on. You have to.

Steve: One of the things that I've had on my mind, Leo, ever since the TrueCrypt discussions, was a project that I hope to take on for the podcast. I think it'd be really

interesting to create a formal definition of security. That is, when we were talking about, well, but these guys are anonymous, okay, well, what does that mean? And like it's open source or closed source. And what does it mean that it's been vetted over time? I mean, it would be possible for us to create a formal definition of what we mean by the word "secure." Where does security come from? Where does it derive?

Leo: Where does security come from, Mommy? Well, son, it all starts when a - actually, by the way, I just want to make one mention. I mentioned that I went to Safari to see if TiVo - and it turns out that preference is per browser. That's a plugin preference to warn for insecure forms. So you do need to go to every browser that you use LastPass.

Steve: Oh, that you - it doesn't go with your account.

Leo: That's why it's not in account settings, it's in preferences. I never knew that. Preferences means per plugin, I guess, or per browser.

Moving on to Question #4 from Kevin Weinman in New Jersey. He wonders about protection from CryptoLocker: I have a small IT shop. Many of my clients use OneDrive or Dropbox for backup. I'm concerned that a malware program like CryptoLocker, once on a client machine, would encrypt those files, and that those changes would be passed up to the cloud. Is there any way to detect the encryption process or change a special permission that would prevent encryption? I suppose one could write a script to open a file and, if it failed, alert the user. But there's no guarantee the selection of files by the ransomware will not be random.

Steve: So we've talked about this before, but I thought it is worth reiterating because CryptoLocker, even though it may currently be sleeping, I don't expect it to stay asleep for long. And we've already seen some clones. I want to remind people of the site BleepingComputer.

Leo: Great site.

Steve: If you just google - yes. If you just google - but BleepingComputer's a little big, and it's even - I couldn't, when I just went to BleepingComputer.com, I couldn't find the CryptoLocker page. So it's better if you google the term "CryptoLocker," and within the first couple links, because the site, as you say, Leo, is so good, you'll find a reference to BleepingComputer. Click that. You'll go to a fabulous page which they've pulled together and are maintaining about CryptoLocker. And Kevin and others, there are things you can do. They're soft fixes. By that I mean they block CryptoLocker today, but the CryptoLocker people have already been evolving CryptoLocker to change some of the earlier advice and recommendations to skirt these things.

So, but to answer Kevin's question directly, on that page, scroll down, you will find a number of settings that actually do completely block it today. So that, if anybody, all of his clients, for example, and you can use a script, or registry settings, or what's the other thing? I can't think of the term.

Leo: Sysinternals?

Steve: Policies. Group...

Leo: Oh, group policy, yeah, group policy editor.

Steve: Group policy system, yes. You're able to use group policies to quickly modify some of the settings on the system that completely block CryptoLocker today. But it is important to understand it may not block it tomorrow because that's the nature of how open our computers are and the fact that the CryptoLocker guys are seeing what people are doing to block them and can probably work around anything we can do. But today it can be completely blocked. Just google "CryptoLocker" and look for the BleepingComputer link.

Leo: And thanks to Web1726, who reminds us that the movie is called "Algorithm."

Steve: Ah, yes. It was originally named "The Hacker Movie," and Jonathan renamed it "Algorithm."

Leo: And it will be screened Saturday night as part of the HOPE Conference. I'm glad to know HOPE's still going on.

Steve: That's the name, HOPE.

Leo: Hackers On Planet Earth. It's a New York City-based conference. I think it's - 2600, I think, does it. And it's a great conference.

Steve: And that's the conference where EFF will be formally announcing their firmware for the routers.

Leo: Oh, good. Oh, good.

Steve: Yup.

Leo: Tom, out in hot Redlands, California - and when he says "hot," I bet it's really hot, like over a hundred - wonders about my refrigerator? Leo mentioned recently that you refrigerate a bunch of your old PDAs. I, like you, have hoarded some Zire 31 models, but I keep them in a regular room, and it can get pretty hot in my area. Why the need for cold storage? Thanks, Mr. G. Cheers, Tom.

Steve: Well, it's chemistry. And you'll remember, Leo, I'm sure, that it was common

practice in the old days that photographers kept their film in the refrigerator.

Leo: Yeah, yeah.

Steve: That's where you put your film. And it's because anything that is involving chemicals is slowed down in the cold. So the reason my PDAs, my old whatever they are, Titanium, I mean, and really...

Leo: I think it's time to really give up on those.

Steve: It really is. I couldn't possibly use them now. But some things, you know, some things don't die, like my HP calculators. I'm glad I have a bunch of them because that's the only calculator I ever want to use. No one has made anything better. It is certainly the case that I would find the Palm whatever it was, Tungsten, no longer able to fulfill my needs.

Leo: Or the Zire 31, yeah.

Steve: But at the time I thought, I mean - yes.

Leo: Hey, but I have a six-pack of Zima in my freezer. I'm going to keep that there.

Steve: Well, and I've got all my batteries in my freezer, all my AAAs and AAs and so forth.

Leo: Batteries are good. Yeah, it's good to do that with batteries, right, yeah.

Steve: For the same reason. You just want...

Leo: But you have to watch condensation. Isn't that a potential issue?

Steve: Yeah. And that's why, naturally, all of the PDAs are in their own little individual bags that are sealed against moisture.

Leo: You need a vacuum machine that you can vacuum seal them so there'll be no moisture in there at all.

Steve: That's right. That's...

Leo: Come on, Steve.

Steve: So the answer to Tom is just because you want to keep them cool to keep them fresh.

Leo: Apparently...

Steve: Film, batteries, and food.

Leo: Except that's not true for lithium ion batteries, apparently. Don't freeze them.

Steve: No, no, no, no. Don't freeze. Absolutely, no, they're not frozen, they're just refrigerated.

Leo: Chilled.

Steve: They're just chilled, yes.

Leo: Chilled.

Steve: Did people freeze their film? Or I think they just refrigerated.

Leo: No, they just refrigerated. In fact, many, many photographers in the film days would have a darkroom fridge or a fridge in their studio where they would keep all their film, yeah.

Kristopher on, he says, "The Internets" worries about microphones being used as bugs: I've always used an external mic with a mixer or physical switches on mics. Yeah, it's normal. Us, too. But when using a webcam, it always bothers me to see signals picked up by the webcam in the recording devices mixer in Windows.

Steve: Yup.

Leo: What I mean is the webcam's mic is only silenced in software. It's muted, not the default recording device, or turned down. But it's there. I believe people with resources, like the NSA or Microsoft, could use these open mics as bugs. Having a mixer, I know from chatting with people online who walk away from a webcam microphone, you can amplify that signal to hear several rooms in a house. Surely compression software could do the same automatically.

I have saved my sanity for years now by going into the Device Manager and

disabling the driver for my webcam mic. Now I have an obsession for an online game where I need to communicate and have both hands. The headset I have has no physical mute button. Is this a concern for the ultra paranoid? Do I have a potential open mic now for the NSA? With a webcam picture, well, you just use tape or point it to the wall. But I don't think people think twice about the mic that's right next to that camera. Steve?

Steve: Yeah. Kristopher is absolutely right. We know that there is software, the various kiddy monitoring software. We've covered stories about schools getting themselves in trouble for loaning laptops to students and then turning the webcams on on the laptops. Certainly the microphones are the same. And we've talked about physically covering up the lens, as Kristopher mentions, in order to block the camera. But the mic is information leakage, too.

The only thing I know is what Christopher has done, which is to go in and essentially remove driver-level, device driver-level support for the device. And in that instance, it just disappears from the inventory that software gets when it says "Give me a list of all the microphones on the machine." But I guess he's using both hands now and using a headset. My advice would be, if he's really worried about this, unplug the headset when you're not using it because otherwise you do have a potential open mic. But, yeah, mics can certainly be used as bugs. And we absolutely know that there is not only commercial software, but we have run across indications of malware, like the RATs, the Remote Access Trojan tools, which among their list of features is "Listen to the room" and "Monitor through the webcam." So hackers know about this, too.

Leo: Probably wouldn't be enough to just uncheck the driver. You'd probably have to uninstall it. I mean, if you've got malware on your system that's turning the mic on, it certainly could reenable a microphone; right?

Steve: I agree. It's not enough to turn it off.

Leo: You've got to get the software off there, yeah.

Steve: And normally, at the API level, where the software works, it has the option to turn it on. So just turning it off doesn't do it. It would see it and be able to enable it and turn the volume all the way up and do everything it needs. Yeah, you're right, you need to actually remove support so that it just disappears from your mixer, and you're not seeing it at all.

Leo: That's kind of a pain if you ever wanted to use it.

Steve: Oh, my god, big pain.

Leo: And don't forget to put your cell phone in a Faraday cage, but of course it has a microphone, too, which is always on. I'm not going to - I don't want to make him

even more paranoid. We've got microphones all over, all the time.

Steve: Yeah, we do. Yeah.

Leo: Steve, that concludes the question-and-answer portion of the show. Is there anything else you'd like to talk about?

Steve: We got it covered. And I think that next week I will have a fascinating podcast prepared, following from that 15-page research report from the UC Berkeley researchers who looked in detail into the operational behavior of those five web-based password managers. That'll be our topic for next week unless something catastrophic happens, and we have something even more important and interesting.

Leo: It is worth revisiting password managers in general because I think we really strongly recommend the use of them. And we want to make sure they're safe and secure to use.

Steve: Yes, you'll need them until SQRL takes over the world.

Leo: Then, never again.

Steve: Never again.

Leo: SQRL, and what's going on with SQRL, available at Steve's site, along with what's up with 6.1 of SpinRite and everything else. He's got so many projects and so many irons in the fire. He also posts 16Kb versions of the audio of the show, plus fully human written transcripts. GRC.com. And while you're there, you might want to pick up a copy of SpinRite, world's best hard drive maintenance and recovery utility. You can get full-quality audio and video of the show at our website, TWiT.tv/sn, or wherever podcasts are aggregated. Just search for TWiT or Security Now!, Stitcher, all of the apps that our wonderful third-party developers put out. In fact, we've now got a list at TWiT.tv/apps of all the different apps and all the different platforms, including Roku, Samsung, Vizio, lots of places. So it's always fun to watch along as we do the show, which is Tuesdays, 1:00 p.m. Pacific, 4:00 p.m. Eastern time, 2000 UTC. Thank you, Steve.

Steve: The Roku app works really well.

Leo: Isn't that nice?

Steve: I was, yeah, I set Jen up with Roku so that I could share some video stuff with her. And I thought, oh, look, there's TWiT is available. And it works great.

Leo: Yeah, isn't that nice? Thanks to Craig Mullaney at ShiftKeySoftware for that one.

Steve: Yes.

Leo: We have a good bunch of third-party developers.

Steve: Okay, my friend.

Leo: Thank you, Steve.

Steve: Talk to you next week.

Leo: Have a great afternoon. Stay cool. And we'll see you next week...

Steve: You, too.

Leo: ...on Security Now!.