News and Feedback

Description: This week we discuss a new bad bug found in the majority of SMTP mailing agents. Fifty-four high-end HP printers are found to be remotely exploitable. More than three-fourths of 433,000 websites are using vulnerable JavaScript libraries. We discuss some horrible free security software, additional welcome Firefox news, a bit of errata, some fun miscellany, and a bunch of feedback from our listeners - including reactions to last week's Quad 9 recommendation.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-639.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-639-lq.mp3

FATHER ROBERT BALLECER: This is Security Now! with Steve Gibson, Episode 639, recorded Tuesday, November 28th, 2017: News and Feedback.

It's time for Security Now!. This is the show with Steve Gibson who reminds all Mac users that we are root. I'm Father Robert Ballecer, the Digital Jesuit, in for Leo Laporte. Steve of course is the big brain behind Gibson Research, ShieldsUP!, SpinRite, and our coming non-passworded overlord. Steve, my friend, it is so good to see you.

Steve Gibson: Well, it's a pleasant surprise. When I heard Leo saying that they were off to the airport a few hours from now, I thought, oh. And I happened to be ready a few hours, I mean, like ready immediately at 1:00. So I sent messages to Leo and Lisa saying, "Hey, I can go now if that would be easier." But then you sat down in front of the camera and began setting up. And I said, oh, it's a Father Robert day. So welcome. Great to have you.

PADRE: It's like a magical wonderland here, and things just happen. There are some very smart people who do scheduling way, way better than I ever could. And they just tell us where to show up, and things tend to work out.

Steve: That's the advantage of having people who are on the job, yup.

PADRE: Indeed. Uh-oh. Steve, of course we're just going to mention it because it just happened a couple of minutes ago, towards the end of MacBreak Weekly, actually, there was a - I'm not even going to call it an exploit. I'm just going to call it someone fat-fingered something when they were releasing the new version of OS X that allows someone to put "root" as the username, no password, and just enter that in a couple of times and then gain access to root for Mac OS X. That's bad. It's going to be patched, I'm betting within the hour or so. But I don't want to cover it because I'm more interested in

figuring out in the post exactly what allowed that to happen.

Steve: Well, okay.

PADRE: Have you heard of this at all?

Steve: So what little we know, because I was learning of it, I was finishing putting the show together when the news hit. So I don't have any additional detail. But there's an important takeaway from our listeners which is that they need to make sure no one has access to their Mac OS X machines in the interim. This is a complete security bypass of all logon authentication for this latest version of Mac OS X. And as you said, you put "root" in as a user, leave the password blank, try it a few times, and I guess it fails initially, and then it succeeds and logs you in with absolute root privilege on any one of these most recently updated Mac OS X's.

So yes, Apple's going to fix it. But what this means - and it's not a remote exploit. But it means that right now there are people, mischievous people learning about this who know they have a small window of opportunity before this gets fixed, during which time they have unrestricted access to any Mac machine that they can find. So any machines that you have physical control over, our listeners who are listening to this, be careful. And you might try it yourself, but don't do anything to hurt yourself if this works.

But anyway, so this is just breaking news. And certainly by the time we know more about it next week, this will have been patched. But right now, certainly this is going like wildfire through the globe. And what it means is, very embarrassingly, at this moment, all of the most recently updated Mac OS X systems are vulnerable to a local attack allowing anyone to get root on the machine and have at the machine. So a complete bypass of logon security.

PADRE: Right. And this is going to work with the High Sierra distro of OS X, so check to see if that's what you've got. And yes, as you mentioned, it's going to be patched almost immediately. I'm sure that there's a lot of egg over at the campus right now on faces. But what has me about this is the fact that it takes a couple of times to take.

Steve: So you agree. It's a weird...

PADRE: That's super strange.

Steve: It's a weird, weird bug.

PADRE: Yeah. This is not a hard-coded password. This sounds like there's a call somewhere that ends up calling to the wrong thing, if you air the authentication [crosstalk].

Steve: If you exceed - exactly. So, for example, we know there are normal lockouts to prevent you from guessing. And/or sometimes various systems will slow down and get more deliberately laggy to prevent you from continually guessing incorrectly. So it makes sense that there could be a multiple failure branch which is engaged if you fail several times. Unfortunately, it then fails open rather than failing closed.

PADRE: Right. Which this is a bit. That's a bit that flipped. Yeah, no biggie. [Crosstalk] in the chatroom points it out. He said, "Look, it's the lockout call." So when it supposed to call for a lockout, instead it calls for this broken authentication, which is awesome.

Steve: A lock in. It's funny, too, because at the end of MacBreak they were talking about both in this context and also about iOS 11. And I disagree with Rene a little bit. He took the position that, oh, well, all major iOS updates, major version changes, have had bugs. I have to take exception. This has been a catastrophe. You can't call iOS 11 anything but an absolute disaster. I mean, I've seen it doing things - it's getting better. And I remember when I saw this thinking to myself, okay, we are going to see the highest rate of updates that we've seen yet. And sure enough, every couple days, in some cases several a day, they're fixing, they're giving us new versions of iOS 11. It's a huge embarrassment. It's just been a catastrophe. I don't want to sugarcoat it.

I'm sorry because I love iOS. I'm an iOS person, all my pads, my phones and so forth. But, boy. And it's rendered my iPhone 6s completely useless. It's just, I mean, I'm managing to use it, but nothing is smooth. Nothing is quick. Nothing works the way it used to. Even just swiping between screens it waits for a while, rearranging icons. It's just broken. So it does make me wish for Jobs being back because this would never have happened under Steve's watch. This is what he prevented from happening. It's why he kept the iPad off the market for multiple attempts, saying, no, it's not ready yet. This is not good enough. And it's just so disappointing.

PADRE: Well, Apple is not supposed to be the "good enough" company. Apple is supposed to be it's polished; it's refined. I mean, yes, you may update every one in a while for new features, but you're not supposed to be updating for simple use. That's weird.

Steve: These devices have always felt like appliances. Yes, there was a computer in there, but we didn't really know about it. We didn't think about it. It was an appliance. It was, you know, we used it for a function. Now they feel like broken computers. They have become broken computers with all of the pain that comes along with something not working the way it should. It's just really sad.

Anyway, this is Episode 639, last podcast of November, for the 28th of November. And the good news is we've been accumulating feedback from our listeners that we've just not had a chance to get to. And for whatever reason, Thanksgiving, maybe the hackers took the week off, I don't know, there just was kind of blessedly not a lot of the typically horrific news that we're dealing with. Well, of course this OS X is pretty horrible. But we essentially have, where we normally have 15 main security topics, we have five today. So I thought, yippee. And so there was feedback from our listeners. And I went back and got the feedback from last week that we haven't had a chance to get to and put it all together. So this one is just called news and feedback.

We have a little bit of news. We're going to talk about a new bad bug found. We've already talked about Mac OS x. Another one found in the majority of SMTP, that's the Simple Mail Transfer protocol. Mailing agents, the one that has I think it's 54% of the Internet uses Exim, E-X-I-M, and a bad problem was found in those, which is like, as in immediately update. Fifty-four, I couldn't believe there were 54, of HP's high-end printers were found to be remotely exploitable by some clever hackers. We'll talk about their work.

More than three quarters of an analysis of 433,000 websites are currently using vulnerable JavaScript libraries. There's some horrible free security software which I actually got clued into from Leo's retweet of an EFF tweet which put me onto that. There's also some additional welcome Firefox news. We talked last week about I think it's 57, which is their Quantum release, which is just running like - it's just screamingly fast. Leo loves it. I've been using it. I'm a fan now. I can't use it on XP, but it's running over on my Win7 machines just fine.

Then we have a little bit of errata, some fun miscellany, and we're going to fill the rest of the show with a bunch of feedback and discussion between us from our listeners, including the first week's reactions to last week's topic, which was the Quad 9 DNS service. I've got a lot of feedback from our listeners after using it for a week. So I think another great podcast to end November with.

PADRE: And actually I'm really interested in the way that we're going to start this because the Picture of the Week is something that was sent to me, as well. Actually, I think the same person copied it to both of us. And actually let's jump into that because it is kind of interesting. This is one of these carbon copy responses that Comcast has been sending out, especially with Net Neutrality in the crosshairs this coming week. They're saying, quote, "We do not and will not block, throttle, or discriminate against lawful content. We will continue to make sure that our policies are clear and transparent for consumers, and we will not change our commitment to these principles," unquote. Now, I have received at least four or five copies of this, Steve. How many did you get from people on Twitter?

Steve: Yeah, same. And so that's Comcast's official tweet that you just read. "We do not and will not block, throttle, discriminate," blah blah blah. And so there is a great response, a tweet from someone whose handle is Lore. And so in that spirit he says: "We never will, but it's very important that we be able to. But we won't. So let us do it. Because we won't do it. Which is why we're spending so much money to make sure we can. But we won't. But let us."

PADRE: And actually that sums it up. And we already know, we know that both Comcast and Verizon have, because it's leaked out of corporate, are preparing to do tiered packages. Once Net Neutrality goes away, they're going to start small. I mean, it's not going to be doom and gloom. They don't want to prove their critics right. So it's going to be simple things, like hey, we noticed that your VPN is going really, really slow. Would you like to buy VPN service from us? They're not going to say we've throttled your VPN service. They're just going to say, if you use our VPN, it will be much faster. It's going to be these little bits and pieces. It's the death of a thousand cuts is essentially what we're going to be heading to, unfortunately.

And the thing is, I don't know if you've gotten into these Twitter battles over the last couple of weeks, Steve, but I've got people who were actually saying, "Well, no, Comcast, they were running ads saying that they support Net Neutrality." I'm like, "Really? So you're going to believe an ad over what they've been doing the last 10 years." That actually kind of disturbs me.

Steve: Yeah.

PADRE: I know about that. All right, Steve. Before we get onto, well, the fun stuff, let's look at the funner stuff. You did find something that I may need to add to my Christmas list.

Steve: Oh, yeah. I ran across it in my Twitter feed as I was catching up on the previous week, as I do every Monday as I'm preparing for the podcast. And it just caught me by surprise. It's just so clever. It's a placemat with a grid pattern which is offered by the Museum of Modern Art, which gives the illusion of it being like a rubber sheet with a grid on it where the knife and fork and the plate are depressing this flexible surface. Or it could also be seen as a gravity well created by the plate and so forth.

Anyway, it's just - I tweeted it. I got a lot of positive feedback from the people following

me who saw the stream and said - because initially I only had a Japanese site. Somehow there's a Japanese version of the MoMA.org site. And so I didn't have a U.S. link. And several people said, "Okay, you just can't tweet that picture without telling us what the URL is." So a lot of interest. And as a consequence I do have in the show notes the link for this and the picture. For what it's worth, I just think it's the best dinner placemat, or I guess breakfast, although it might be a little much for breakfast, that has ever been conceived. So I just loved it and wanted to share it with the podcast listeners as a completely off-topic but very fun thing for - and as you said, maybe perfect for the holiday season.

PADRE: You know what I love about that? The fact that you could put these down at a table, bring over your friends and family, and just look around. And you'll immediately know which ones of your friends and family are geeks, like real geeks, because they'll go, oh, time-space displacement. And the others will just go, oh, it's kind of a funny mat.

Steve: Yeah. So I mentioned before that a very worrisome remote code execution and DDoS, actually two different exploits, two different vulnerabilities, were discovered in the Internet's majority Mail Transfer Agent, MTA, which is the official term for these things. These are SMTP servers, Simple Mail Transfer Protocol, which forward mail around the Internet. Exim, E-X-I-M - I don't know if that's how you pronounce it. Is it xim or ex - I've always just said Exim.

PADRE: I've always said Exim.

Steve: Yeah. So it is 54% of the MTAs, the Mail Transfer Agents on the 'Net, are using Exim. It was designed for Unix-like operating systems, and like Send Mail is the even older venerable sort of original Unix implementation for mail. Exim is easier to use. It's available on most Unix-like systems. It's been transported, or it's been ported, to Microsoft Windows using Cygwin. And it's the default MTA on Debian and GNU Linux systems. So it's around. It's very popular in the U.K., where ISPs and universities are using it. And also it's widely used with GNU's Mailman mailing list manager and also the widely used cPanel. Oh, it's 56% rather than 54. I remembered the number wrong.

Anyway, so against this backdrop of this thing being widely used, over the Thanksgiving holiday a security researcher discovered and publicly disclosed two critical vulnerabilities in this majority platform which is by its nature public. I mean, this is public mail routing. So they're all publicly accessible. There are two vulnerabilities. The first one is a class of vulnerability we've discussed in the past, a use-after-free bug where some memory is allocated. Then it is released back to the operating system, yet the code still has a pointer to the released memory. And there are ways then that it's possible to leverage the use of that after it's been freed in order to get code to be executed from that freed memory.

And since this can be supplied remotely and as a consequence of this vulnerability, it can be loaded with code that the remote attacker provides. It creates the ability to remotely inject malicious code into the server by - in this case it's a special sequence of BDAT commands, which is a class of command that the MTA uses. So that's the first one, basically publicly vulnerable, remotely injectable, remote code execution in a vast number of machines. A Shodan scan revealed that more than 400,000 servers are currently vulnerable.

So the takeaway for every one of our listeners is, if you know of or are responsible for an Exim server, you want to update this immediately. The updated version 4.9 is now available on GitHub. So it's there. 4.88 and 4.89 are the vulnerable versions. The second vulnerability of the two is a DoS, a Denial of Service flaw which essentially prevents the

server from accepting and processing SMTP mail connections. So all it does is it just locks up that server and prevents it from performing its function, thus denying the service that it was designed to provide. But probably, given that it is scannable, Shodan provides them.

There are 400,000 of them at this point. Until they are updated, anybody within range of this audio absolutely wants to update as quickly as they can. Oh, and a Python-based proof-of-concept code has been released, so there's even a template for how to do this. The code itself is not malicious. It's a proof of concept. It just puts the server into a five-second loop in order to demonstrate that remotely injectable code is successfully running in the system.

But everybody wants to - anyone who has access to or responsibility for one of these wants to update to 4.9 immediately because we can find them, they're public, proof of concept makes it easy to implement, and in no time we're going to see exploits of this, essentially people taking these servers over and who knows what, botnets or command-and-control waypoints, who knows. So not good. And the link is GitHub.com/exim/exim in order to get the latest update. I've got the link in the show notes.

PADRE: Now, that second one, the second CVE, that's limited usability because essentially it's a DoS attack. And I thought maybe, if there was a way to cut all connections except the connections that a remote attacker was using, it could be usable. It'd be nice because you could cut off all other activity. But it would cut even the connection that the remote attacker is using. So that's essentially a nuisance attack.

Steve: Right.

PADRE: But that first one, that is kind of scary because that's built in. Now, run me through this. So the BDAT, it's running within SMTP. I make it generate an error. And when it generates an error, actually it carves out an allocation of memory, and that's the memory I use to try to run a remote code execution?

Steve: I didn't dig into it because it's already been fixed, and the bad guys are going to be cleaning this up or taking advantage of this. And it has been fixed. There was some comment about ending a command with a period, but I think that was the second, the DDoS flaw. And so from the CVE it says that you can remotely execute arbitrary code in the SMTP server by crafting a sequence of BDAT commands. And again, there was no point at this point, for me anyway, digging any deeper. The bad guys are certainly going to want to. And all the information is there, especially in the Python-based proof of concept, for how to take this thing all the way to the mat.

PADRE: Oh, okay. Actually, I think I remember reading about this. The issue is that it's assuming that, if you're putting any code into a memory allocation, that you've compiled it with Pi because Pi will keep you from doing anything malicious with it.

Steve: Correct. I did see that, yes.

PADRE: But an attacker wouldn't be doing that because an attacker wants to do something malicious.

Steve: Exactly.

PADRE: So it makes the allocation, and then they can just shuttle in their code.

Steve: And put anything they want, exactly. And then release it, and it gets used.

PADRE: Well, that sounds fun.

Steve: Yikes. Yeah. Yeah.

PADRE: Now, I could see someone doing something like doing the first exploit and then doing the second exploit to cover their tracks. They'd hang the machine. Someone would just figure that it's gone into some weird state, and they do a restart. And now everything is running the way that the attacker wants it to run.

Steve: Well, and what we see, especially for something like this, is if there are 400,000 MTA machines, half of them are virtually abandoned administratively. They're in closets. They're dusty. No one has looked in on them for years. They're not going to get fixed. So there are certainly, you know, the ones that ISPs are running, hopefully ISPs are on the ball. Well, universities are probably going to be the hosts of some of those that are in the closet that have been long since forgotten. These are just workhorse servers that nobody thinks about because there's nothing sexy about shuttling mail around. And especially if it's just a server that accepts mail and forwards it. It doesn't even have like a body of accounts that clients are connecting to, necessarily. It's just a forwarding agent.

So these things tend to not be visited often. And I think it's entirely foreseeable that, as a consequence of this, we now have a new inventory which, for the foreseeable future, bad guys are going to be able to leverage into obtaining presence in as a consequence of this flaw.

PADRE: Wow. We'll see. And fortunately, yes, there is a patch. And most of those machines will never be patched. in fact, I would leverage a guess here that the majority of those machines haven't been patched since installation.

Steve: Precisely, yes.

PADRE: Well, that's some good news.

Steve: Well, actually, if the problem was introduced in 488 and 489, then older ones may not be vulnerable. This may be a relatively recent update which introduced the problem at 489. And so maybe, as a happy coincidence of this, it's the ones that have been maintained that now need just a little bit more maintenance.

PADRE: Isn't that a kick in the pants.

Steve: Yeah.

PADRE: So that's your punishment for updating your servers.

Steve: That's right.

PADRE: Well done.

Steve: So I didn't realize that - and I guess this must be history also that HP has 54 models of enterprise printers. I mean, that even exist, 54. Anyway, it turns out that HP in some of their marketing material was annoyingly boastful, talking about how bad printer security is of everybody else, but they're wonderful.

And in fact on the site, I didn't bother to dig into this video because it wasn't germane. But on the site of the security guys who decided to take a closer look at HP printers, they have this video which, from their description, which I did read, it just sounds obnoxious. It sounds like HP marketing getting a little carried away with themselves, thumbing their nose at the poor security of everybody else's printers except HP. So these guys said, "Ah, okay. Let's take a closer look at HP printers." So you're going to get a kick out of this, Father Robert.

So first of all, years ago, I remember us talking on the podcast about how enterprise-class printers contained hard drives. So my first reaction even then was, what? What? A printer has a hard drive? So and back then the worry was that decommissioned printers that were sent off, being recycled or being thrown away, essentially, had hard drives on which were records of maybe every print job it had ever done. I mean, so obviously a huge privacy and security concern because these hard drives were caching print jobs and leaving them there, just because the drives were not being wiped. Nobody was aware that there was gigs of nonvolatile storage that had been saving the print jobs in a printer. Why would a printer have a hard drive?

Well, turns out they do. And these HP enterprise-class printers also have hard drives. Now, the first thing - this is so fun. The first thing these guys noted is that, wow, okay, HP apparently wants people to not poke around in the operating system and firmware of these devices. So the printers use FIPS standard, F-I-P-S standard-encrypted, hardware-encrypted drives. So, and we've talked about this technology. When the drive boots, a password is given to it through its interface to unlock a private key through which a symmetric cipher decrypts the contents of the drive on the fly as it comes through the electrical interface so that the system gets decrypted hard drive content, but the drive content is encrypted. And unless you have the key, you're not able to see the hard drive.

So they got a printer. They purchased two for this purpose. They take the drive out and mount it, and it's encrypted. So, okay, they go, whoops, okay, it's encrypted. What to do? So they're clever. They got a non-encrypted, that is, non-encryptable drive, some other drive, I think it was a Toshiba that didn't offer hardware encryption, stuck it into the printer. Now the printer says, "Oh, I have a problem. I don't have my OS and firmware. I can't read it. We must have had a hard drive problem." So they use a thumb drive to reload the OS and firmware into the printer. Which is now not encrypted on the hard drive because the hard drive doesn't support encryption. Now they take the drive out and mount it and have access to the file system.

PADRE: Oh, so HP was using the built-in hardware encryption inside of FIPS-enabled drives. So you just put in a non-FIPS-enabled drive, and it just says, okay, I guess this is all right?

Steve: Okay, yeah.

PADRE: That sounds like a bad design, Steve.

Steve: So, yes, they download the OS and firmware in the "restore your broken printer" mode, onto a drive that can't be encrypted. And now it's not.

PADRE: It sounds like that's something you would check for. Actually, it does check for it, and then it just says, eh, okay.

Steve: Yeah. We gave the drive an unlock password that it ignored, and everything's fine. So then they discover that - and this is where it starts getting, like, really? It's

running Windows, of course, because of course. It's running Windows CE. In their link - and I'm not seeing the link here. Oh, yeah, there it is. You might want to bring it up. It's FoxGloveSecurity.com. The link is an amazing walkthrough of what they did. But they show the Windows directory with all the familiar Windows files. DLLs, there's ATAPI. dll and all the file system files, I mean, it's very familiar to anyone who's ever poked around in Windows. So yes, indeed, it's running Windows like all of our favorite kiosks around the world.

So they then proceed, now that they have the OS in firmware, to reverse engineer it. And I won't go into the painful details. But essentially what they determined was that both HP solutions and firmware updates, HP has this thing called HP Solutions which allow an extension of the functionality of the printer with a proprietary SDK which is not available, but is made available to major HP partners to allow them to create their own updates, their own additional features packs, which are then signed by HP in order to authenticate them.

So they write: "Both HP Solutions and firmware updates consist of a single file with the .BDL, short for bundle, extension. This is a proprietary binary format with no publicly available documentation." They write: "We decided that reverse engineering this file format would be beneficial as it would allow us to gain insight into exactly what firmware updates and software solutions are composed of."

So what I'm not going to go through in detail - it's all there on their site for anyone who's interested - is the blow-by-blow of their reversing of the file format. They find CRC32. They make some changes. It doesn't work. They dig deeper. They find other things. They end up completely reverse-engineering the structure of the file, so much so that at GitHub.com/foxglovesec/hpwn, which is to say H-P-W-N, they offer a Python script called BDL Patcher. Okay, this is the proprietary format, signed cryptographically, but not so much, signed bundle.

And the description for BDL Patcher says: "This Python script" - freely downloadable - "can be used to create modified, but still valid, HP software solution bundles. Usage instructions are built into" - how convenient - "built into the tool. All you need to do is open mod.zip, replace or add any files that you would like, and then run the following command: Python hp_solution_patcher.py [blah blah blah]. This will generate a new, modified BDL file called patched.bdl." Which the HP printers will then happily ingest, they will accept, and allow you to run your own code on the printer.

PADRE: This is so weird.

Steve: I know. It's just like, oh, my god.

PADRE: Wait. But Steve, how can you make a modified and yet still valid BDL? That makes no sense to me, what kind of checking it's doing, because obviously it's not doing any sort of checksum. It's just the certificate? Just a signature?

Steve: Apparently HP, unfortunately, rolled their own format. They said, oh, let's put some checksums here and some here and some here. And this will be really hard, and nobody will be able to figure it out. And besides, we're encrypting our firmware on this hard drive so nobody will ever get to see it anyway. So, whoops.

So they say: "There exist a number of methods for updating the firmware of HP printers. Most administrators would be aware that firmware updates can be installed through the printer's web interface or through the Web Jet Admin client." They write: "Firmware can

also be installed at boot time through BOOTP/TFTP [the Trivial File Transfer Protocol] options. Additionally, the Security settings page on the HP printers implies that firmware can be installed through a print job..."

PADRE: Oh, good, yeah.

Steve: Here, take this, HP, "...over port 9100." Yup, the standard HP printer port. So they write: "Thus there are a number of avenues available for remote network injection of malicious firmware into printers." So there is a list of the printers, linked to from their report and also here in the show notes, of the 54 printers which will accept, basically, any code that's - it's Windows, Windows CE. So add your malware, wrap it in this .BDL format, send that printer a job it will never forget. And it will write it onto the hard drive, and you can take the printer over.

PADRE: And the list is pretty extensive. It is the most popular HP enterprise printers. It's the M651, the M680, the M631, and the X556. Those are their top sellers.

Steve: It's the good ones, yes.

PADRE: Those are the good ones. And the crazy thing about this is, well, two things. First, it really sounds like HP was putting all their security eggs in the encryption basket. They just figured, if no one can see this, then they'll never be able to backwards engineer it. So this was a clear case of security through obscurity. They kind of threw a nod by doing some checksums, but they were clearly just hoping that no one would ever figure out a way to actually get the files.

And the second thing is, because you can replace the firmware, and because this is essentially a Windows computer, this becomes a pivot point. I compromise one printer in an enterprise. That printer is trusted by the entire network because everyone needs to print to it. And now I can probe that network as if I had a Windows machine inside the network because that's what I have. That's scary.

Steve: Actually, we have in the past encountered instances of APTs, Advanced Persistent Threats, living in compromised printers. I mean, that's why printer security is a thing HP is bragging about, denigrating everybody else's security, saying how fabulous theirs is. Which in fact is what induced these guys to take a closer look. And when they did, oopsie.

So HP has responded. There is an update to the firmware for all these printers. And once again, anyone in an administrative capacity who has access to or has responsibility for any of these printers should seriously consider updating the firmware. Unfortunately, normally printers aren't doing that themselves. You've got to log into the web interface, check for updated version. The printer will say, oh, what do you know? I've got new firmware. Then you've got to install it and restart the printer and so forth. And as you said, Father, what they should have done, but didn't, was to verify that the ATAPI feature set being offered by the drive includes encryption, and to refuse to put that drive online in any way if the drive it encounters does not support encryption.

PADRE: Well, that makes sense. If that was going to be your security deterrent, then you need to make sure that it's not as easy to bypass as, oh, I'll just install a different drive.

Steve: And I'm sure they're doing that now with this update. They're like, oopsie.

PADRE: That sounds like a flag that should be set: If no encryption, just don't boot.

Steve: Well, yes. And really there is no excuse for not using a cryptographic digital signature around this bundle package, so that the only thing that exists in hardware is the public key, which is used to validate the private key, or the signature made by a private key that only HP has. They didn't do that. They just said, oh, let's do some CRC 32s. It's like, what?

PADRE: Yeah. I'm trying to research right now because it may have changed, but when I was still doing active administration, the HP printers that we purchased had two firmwares. And it was this whole idea, if one gets corrupt, it will just copy over from the backup firmware. If they still do that, this could be a persistent threat. Unless you actually overwrite both sets of firmware, then the compromised firmware can actually still exist in memory. I'll have to look that up. Maybe they don't do it anymore.

Steve: Normally you do want to have a fallback in case an update crashes.

PADRE: Right, exactly. That's what it was designed for. So if you're pushing an update over the network, and it's incomplete, it doesn't just brick the printer. It says, okay, I had a bad installation. I'm going to copy over the original. So, well, now I've got homework, Steve. Great. I think we have one of these at the school. I could take it offline for a few hours.

Steve: Yeah. It would be a good thing because, again, it's not clear whether these ports are publicly exposed, that is, on the public Internet. This is probably largely an Intranet problem. But we've seen many problems with configuration where it was like, wait, your printer's port 9100 is mapped through to the public Internet so that somebody can, oh, yeah, well, we want our telecommuters to be able to access it, blah blah blah.

PADRE: Precisely.

Steve: So you would just want to make sure that it's, well, I mean, you certainly don't want even Intranet because, if something gets in, then this allows it to establish a foothold in a place that no one would think to look.

PADRE: Yeah. I mean, if I were trying to actively exploit this, I would just be running nmap sweeps right now against port 9100, just looking for 9100. And anywhere I find 9100, I'd enumerate it. If it's an HP printer, I've just scored.

Steve: Yup.

PADRE: Well, that's okay. That's good news, then. So well done, HP.

Steve: Send it a print job it'll never forget.

PADRE: The last print you'll ever need.

Steve: So a recent study of open source JavaScript libraries in use on websites produced some disturbing results. They looked at 433,000 sites and discovered that just over three quarters of them, 77% of the 433,000 sites checked, were currently using vulnerable JavaScript libraries, that is, libraries known to have at least one vulnerability. And they also found that where sites had one, they more often than not had more than one. So not surprisingly, jQuery was far and away the most often seen library, with 82.4% presence in the 433,000 sites. So jQuery at 82.4, jQuery UI at 19.9, Modernizer at 15.1%, Bootstrap at 13.7, and then downwards with popular ones. I've enumerated them here in the show notes.

So that's not vulnerability, that's the adoption percentage. So jQuery far and away, at 82.4, the leading JavaScript library that is present on these 433,000 sites. The number of times it was found vulnerable was 92.5%.

PADRE: Wow.

Steve: Yes. The numbers are high: jQuery UI, 89.7; Moment.js, 73.0. Even though Moment was way far down the list, it only appeared in 3.4% of the sites, it was still three quarters of the time vulnerable. So high level of jQuery adoption. High-level of jQuery vulnerability, which is surprising. Then we look at, okay, wait a minute. What's going on? The oldest version of jQuery with no known vulnerabilities is 3.0.0, released in June of 2016. So what we're seeing is that, once sites go production live, they stop staying current. They stop updating the libraries that they're using, and they just stick with what they've got.

PADRE: I mean, that's human nature. Once it's working they'd just as soon leave it alone because you might break it.

Steve: If it's not broke, don't fix it. Unfortunately, and what isn't clear, I mean, I don't mean to sound like the sky is falling here. When we say "vulnerability," they could very well be minor things. I mean, it's not - if the sky were falling on 90-plus percent of websites, it would have been fixed. The sky would have been patched and fixed. In this case they're known problems. But what's interesting is that, recently, everything is fine. That is, the current libraries are running known vulnerability free. And so if these sites were current as of the current version of these packages, they'd be fine. It's just that they are at least about a year and a half out of date. And those older packages were vulnerable, and they've not been fixed since.

So they write in their coverage: "Each of the front-end libraries most commonly found to be vulnerable has been free of known vulnerabilities for anywhere from one to five years." Which is to say that this massive install base of open source - largely jQuery, 82.4% - 92.5% of those are older than a year and a half and have vulnerabilities that have long since been fixed, and those sites have not been updated.

PADRE: Right. I'm looking through the list of the vulnerabilities. And you're right, most of them are minor. They can either cause a system lockup, or they'll just cause a slightly unexpected return. But these little vulnerabilities, when you pile them up, and especially when they're this old, this is how you develop an exploit.

Steve: Right.

PADRE: You combine a couple of these, and suddenly you've got remote code execution. So that's - and the fact that they've been out there for so long, and you can essentially count on the fact that most of these sites will never be patched, that's, again, then we get into sort of ho-hum worrying territory.

Steve: Yes, yes.

PADRE: Hmm. Now, it's interesting because, if you look at the list, Google Maps is on this. Wait, what?

Steve: Boy.

PADRE: The Google Maps library?

Steve: Yeah.

PADRE: Okay. Let's get back to the action. So we've got vulnerabilities everywhere that are piling up, and people aren't going to patch. But I'm more interested in finding what the Treasury Department is doing in terms of fraud investigations because this one's interesting, Steve.

Steve: Yeah. So three years ago the EFF posted a detailed takedown of a really disturbing piece of software known as ComputerCOP, Computer C-O-P. And for years it was being offered by a company somewhere in the Northeast, I can't remember where, like New Jersey or New York somewhere. It was being purchased in bulk by law enforcement agencies around the U.S. and then freely distributed by them as a public service and as a public relations freebie to their communities. So, for example, in one case a Sheriff's Department bought a copy for every family in its county. And the idea was...

PADRE: [Crosstalk].

Steve: Yeah, well, yeah, huh. So, and at the time of their reporting, the EFF found that this ComputerCOP software was being distributed by 245 different agencies in more than 35 states in the U.S. Okay. So what's the problem with ComputerCOP? Well, among its several features, first of all, it's just - it's not great software. It's the kind of thing that Mom and Dad install on their computer that proposes to allow them to find out what images their children are using, except it turns out that it only knows where the images of IE and Safari are located, so it doesn't find any images if they're using any other browser. And it does no filtering of them, showing all of the icons and favs and all the little Like buttons, I mean, it's a very, very poor piece of software - poorly written, poorly designed, basically just a scam.

So, I mean, it sort of does what it purports to, but it's not anything anyone would want to use, so it's worth what you pay for it, or what Mom and Dad pay for it, which is nothing because it's just being given out by law enforcement, being called the first line of defense against the evil, the dark web, the problems on the Internet.

Well, among its features is a keystroke logger, which records the keyboard activity for all of the computer's users; stores the entire keystroke archive, unencrypted, on the systems hard drive, meaning credit card numbers entered, passwords entered for all the sites everyone on the computer visits. Account names, everything. Every keystroke. When it encounters a keyword which might be questionable, a trigger word, it emails the unencrypted - unencrypted - stored archive to a central server, which then emails it back to the household's parents.

PADRE: That's such a good idea.

Steve: All in the clear. All in the clear.

PADRE: Okay. So let's break this down. This is essentially - it's not quite a RAT, it's not a Remote Access Tool, but it's pretty dang close.

Steve: It's an auto RAT.

PADRE: Yeah, it's an auto RAT.

Steve: It's an auto RAT.

PADRE: So it sits there, and it keylogs everything.

Steve: Yes.

PADRE: And then upon receiving...

Steve: Because you never know, you never know what nasty words someone might type.

PADRE: Precisely, precisely. And then upon the receipt of some sort of trigger word which says that I want a list of everything, it doesn't just provide it to you. It mails it, unprotected, unencrypted, in cleartext...

Steve: To a third party.

PADRE: ...to a third party which then mails it back to me.

Steve: Yes.

PADRE: In unencrypted, unprotected format.

Steve: Yes, yes. There's email containing...

PADRE: That's a good system.

Steve: ...all the usernames and passwords and credit card numbers and everything that's been entered into the computer's keyboard.

PADRE: So essentially anyone listening to the line would just receive a cleartext stream of everything that you've typed in your computer.

Steve: Exactly. Anybody doing a packet capture, anyone watching the Internet is going to see, in cleartext, all of the private keyboard activity on that machine. So to increase the tastiness of this for law enforcement, it was at the time that the EFF brought it to the world's attention being marketed with invalid endorsements, purportedly from the U.S. Department of Treasury, which has since issued a fraud alert over the document, which the publisher of ComputerCOP doctored. ComputerCOP's publisher claims an apparently nonexistent endorsement by the American Civil Liberties Union, the ACLU, and an expired endorsement from the National Center for Missing and Exploited Children. Even if those endorsements were legitimate, they're not backed by any responsible investigation into the design and operation of the product.

That was three years ago. Today, the EFF is able to report that the U.S. Treasury Department, speedy folks that they are, have concluded their fraud investigation into this ComputerCOP, quote, "Internet safety" software. The Treasury Department did find that the company had forged documents which it had been using as marketing material for its products, and that those forged documents were instrumental in the sales of the software to U.S. law enforcement agencies across the country.

PADRE: That's where you get in trouble, yeah.

Steve: However, the three-year statute of limitations had since expired. And the forged documents were no longer present on the company's website or marketing materials. So there's no penalty that can be brought against these cretins who are publishing this. Even today, the software continues to be sold and offered free of charge to the public where the software is promoted to the law enforcement agencies as "perfect election and fundraising tools."

PADRE: Well, there's different levels of meaning for that because, if you want to spy on people, yeah, that's a great tool. The amazing thing about this, well, a few things. One is that they would use the ACLU as one of their fraudulent endorsements. I mean, the ACLU is basically a pile of lawyers. Why would you ever do that? That's just stupid for self-preservation.

The second thing is, and this bothers me, Steve, yes, the statute of limitations is over. But they were going to get in trouble, not for the fact that they created horrible, horrible software, but for the fact that they marketed it wrong.

Steve: Correct. Yes, I know. In the EFF's follow-up just a day ago, they wrote: "In 2017" - that is, this year - "the Lake County Sheriff's Office in Florida purchased 1,000 copies for $5,975" - by the way, this is using public funds. This is funds from evidence recoup slush fund or something. So nearly $6,000 was spent for a thousand copies of this crapware, according to SmartProcure, which is an oversight database. And McGruff the Crime Dog was handing out copies...

PADRE: Oh, not McGruff.

Steve: McGruff has been subverted, "...handing out copies" - in his hot little dog suit - "this summer at a community screening of the film 'Elf' in Islip, New York."

PADRE: They dragged McGruff into this, Steve. Now, that's fighting words.

Steve: You can't make this up, Padre. You can't make this up.

PADRE: I actually kind of want a copy of it now. I just want to see [crosstalk].

Steve: Ooh, I'm sure you can find one. I bet you can find one. The takeaway for our listeners, of course, and for anyone they know, is to stay as far away from this ComputerCOP spyware as possible. What a catastrophe. So thank you to the EFF for, you know. And no thanks to the Treasury Department for not moving quickly enough to slap these people with a fine that would knock them out of existence forever because this crap should not be provided anywhere.

PADRE: There's something larger here, Steve, and I actually saw it over the break because I went to be with my family, my parents, in Las Vegas. And I basically had to nuke every machine in the house, aside from the Chromebooks that I had given to my Mom and my Dad. And it's because my father kept downloading software. And he's like, "But I paid for this. I paid for it." I'm like, "Dad, just because you pay 10 bucks for a piece of software doesn't mean that it's good." And there was every kind of malware installed on these things. And this is not a singular thing. This is a particularly bad example. But there's software out there that, even if it isn't straight up malware, is just so poorly coded that it doesn't belong on a computer of anyone that you like.

Steve: Yeah.

PADRE: Well, I think it's back to the "I'm locking down your computer so you can't install anything without me."

Steve: Yeah.

PADRE: Oh. I don't want to do that.

Steve: Okay. So...

PADRE: Give me some good news, Steve.

Steve: One happy piece of good news.

PADRE: There we go.

Steve: Then we get on with quicker bits. And that is something interesting. Firefox is planning to flag sites that have been hacked in the past, ever. Working with Troy Hunt, who we often speak of - he has the HaveIBeenPwned site and database, which currently has a list of 254 websites that have had their users' password databases maliciously exfiltrated in the past. A new feature planned for Firefox will check any sites Firefox's users visit and preemptively warn them when they're visiting any site that has in the past suffered a breach of its users' secret data. So, for example, you go to Yahoo.com.

PADRE: [Buzzer sound]

Steve: There'll be some sort of a warning, just telling you, oh, by the way, this website has had its data lost. And naturally breached websites won't like this. But users will. And the news of this has been greeted with some happiness on the part of users. And of course we could hope that the specter of this kind of long-term persistent shaming may serve to keep companies more focused upon not being added to that list in the future. It's not clear whether there will be a date or a number of - like what additional information will be made available, what sort of metadata might be included. But Troy has confirmed that he's working with Mozilla, and they're in discussions about how Firefox can pull from his database to keep itself current and just let people know if they visit a site who at some point in the past has not been sufficiently secure.

So it's interesting. We're seeing, as web browsers become more centric to everyone's experience, I mean, they're now the application platform that many people use with Office 365 and everything being - and all of Google's stuff being cloud-based. We're seeing browsers being increasingly proactive now in coming up with useful security and privacy features for their users, which I think is just all for the best.

PADRE: The thing about this, though, is I have seen a complacency among users when they get the warning window, be it whatever browser it is, that this site is not secure, it's not using HTTPS. And they've gotten desensitized to it, so they just click through, click through, click through.

Steve: Yup.

PADRE: They need to do something different for this because this is not a warning in that sense. It's not saying, hey.

Steve: And it does not stop you from having access. It's just an advisory. Probably a little shim up at the top across the top of the page that we sometimes see browsers

offering, just saying, oh, by the way, you should know.

PADRE: I think that needs to be something that, if you're a security professional today, and you design products, you need to start looking at UI a bit more because you can give all the great information to the user, but if they don't actually understand what you're trying to do, they see it as just a nag screen.

Steve: Right.

PADRE: And they'll click through it just like they do when they click through EULAs upon installing software. But there has to be a little bit of thought into, okay, here's a behavior that I want to encourage among my users. What's the best way to encourage that? And it's not just throw something up onscreen that I can click through.

Steve: Right.

PADRE: If someone goes to the Equifax website, there needs to be something that's very quick and to the point that says "You may not want to trust this entity." Because you mentioned that maybe it's time for these companies to feel a little bit of shame. But I think you also have to do the opposite. There has to be sort of a carrot-and-stick approach. You have to be able to hold up the companies that do it right.

Steve: Yeah.

PADRE: Well, okay. You know what, that is a little bit of good news, Steve, so thank you. I do appreciate this.

Steve: Yes, yes. So we've got some errata. I wanted to thank all of our listeners. So many people reported that I had, on the Security Now! page at GRC.com, when I started posting the November updates, so November, what was it, 7, 14, 21, and now 28, I forgot to change October to November. So for everybody who said, "Steve, you got the month wrong for the last three weeks," thank you. The month is now correct.

PADRE: You know what, it is always the same month in Gibson world.

Steve: That's right. Things move slowly here, but we try to keep things clean or corrected. Speaking of corrected, I made a mistake last week that I wanted to correct, thanks to someone posting as CHY, just sent me a short note saying: "@SGgrc Correction about the role of USCENTCOM, Steve." He said: "It's not inherently an 'intel gathering op,'" as I had said last week. He wrote: "It's a theater-level combatant command under which there are many, many sections." So just to correct the record, thank you for that.

PADRE: Yeah, it's always nice to know the internal organization of that group.

Steve: Yeah. In the past we've talked about, and they've been very popular, Humble Bundles, which are these pay what you want for collections of eBooks of various sorts. The most recent one was a security suite which had among them some outstanding security texts. I know many of our listeners jumped on that.

There are currently six days remaining in a Java language-centric Humble Bundle which looks very good. It's basically all of O'Reilly's Java books. And, I mean, just it's a phenomenal - if Java is a language which impacts you, and in fact I'm surprised at how successful Java is, especially in the enterprise, it's because it's based on a VM that is

then able to run on cross-platform. It's a write once, run anywhere.

I mean, it's taken a bunch of hits from us on this podcast in the past because it was a source of many - the Java Virtual Machine was a catastrophe from a security standpoint, and Java itself should have never been given a web-facing presence. But as an implementation language for non-web browser solutions, I think it's the leading language for enterprise implementation. Anyway, this Humble Bundle is very inexpensive. It's just a broad sweep of O'Reilly books. I've got the link in the show notes. Six days remaining as of today.

The second is a science fiction Humble Bundle that has eight days remaining. And again, a huge number of eBooks available. I saw one author who I'm a fan of, Reynolds, who writes sort of dark, creepy sci-fi; but I've read a couple of his, and they're just kind of interesting, but good. Alistair Reynolds is the guy I'm thinking of. Anyway, so for what it's worth, if you just like - I think you've got it pulled up there, Padre. So it's like, what, for $15 or something, it's just a big collection of sci-fi. If someone just wants a big input of new eBooks, I wanted to bring it to our listeners' attention.

PADRE: No, I love the Humble Bundles because I tend to stay to the same sort of themes of sci-fi. Like right now I've been big into the Monoverse.

Steve: Okay.

PADRE: And making a jump to the next and the next and the next is not always comfortable for me. In fact, I'm a little weird in that I will read and listen to the same book over and over and over and over again.

Steve: Well, in that case join me because I'm currently rereading my favorite series, which is the Frontiers Saga, which is now the first set was five, and the sixth book is now published in the second set of - I'm sorry. The first set was 15. And now we're six books into the second series. And so I'm rereading - I immediately started rereading them all. They are so good. So, yeah, I'm with you. I reread books I really enjoy just because I love the experience of suspending my disbelief and getting completely absorbed.

PADRE: I do the Whispersync for the Amazon ecosystem. And I've got to - I love it. I did that with the Expanse series. And I've probably listened to the entire Expanse series at least five times, maybe even six times. I just keep going through.

Steve: Yup. And of course I read all the books, and they are fun. And the series on Syfy, in fact we do have a question coming up in our closing the loop is about that. But I found just - we've been talking recently about bitcoins. And I have got a good friend, Mark Thompson, who is a miner of bitcoin and who is in fact considering moving to a different state because he's found a huge warehouse and is planning to invest a million dollars in mining hardware in order to mine. That's something he's doing.

But what's happening with this phenomenon is interesting. There's a U.K. site, PowerCompare.co.uk/bitcoin, that has pulled some statistics together. For example, bitcoin mining in aggregate is now consuming more electricity than 159 countries, including Ireland and most countries in Africa. So, yes, admittedly, they're not massive countries. But still, I mean, bitcoin mining power consumption is a thing. And just to give people some other little bullet points, in the past month alone bitcoin mining electricity consumption is estimated to have increased by 30%. In the past month it has grown by 30%.

PADRE: That's unsustainable.

Steve: Well, yes. And if it keeps increasing at this rate, bitcoin mining will consume all of the world's electricity by

February of 2020.

PADRE: Oh, good, yeah.

Steve: Which is when we become a black hole and gravity exceeds our ability to pull away.

PADRE: Who knows? I mean, maybe this will spark a new gold rush for power generation. Someone will finally come up with some very cheap and very plentiful power so they can mine.

Steve: Well, the estimated annualized global mining revenue is now at $7.2 billion U.S. Whereas estimated global mining costs is at 1.5 billion. Which means there's some profit in that.

PADRE: Yeah, there's still money to be made.

Steve: There's money to be made if you're the person driving the coin up the curve.

PADRE: But the money to be made, though, is not in the mining. It really isn't, not anymore.

Steve: Right.

PADRE: The money to be made is just having the bitcoin and watching the price go up.

Steve: Yes, well, yes. Just incredible. Incredible.

PADRE: We've touched on this before, though, Steve. Because, I mean, there's actually a cost to running the blockchain that is not associated with mining. Mining is a huge part of it. But it actually does cost energy for every transaction that you do in the blockchain.

Steve: Yes, yes.

PADRE: And you could say the same thing about the Internet. A lot of people will look at their phone, and they'll say, "Oh, this is a very low-power device." Yeah, but everything that it's connected to to make it useful, that uses up a lot of power for all they'd capture.

Steve: Right, right.

PADRE: I don't know.

Steve: So I didn't know I was going to have you here when I had a tweet from a Chad Emory, who tweeted to both you and me. He said "@padresj @SGgrc," he wrote, and this was - I just ran across this in my Twitter stream in the last, well, yesterday when I was pulling things together for the podcast. He says: "I keep hearing of people using SpinRite on Android phones." Of course you famously did, Padre.

PADRE: Right.

Steve: And he said: "Standard build and MS boot build will not see phone on USB. Is there a trick on how to get it to work with Samsung phones?" And I didn't know I was going to have the benefit of your experience, which you've shared with us, about using SpinRite to successfully recover a phone. But I did want to tell Chad that the trick is that until - and I'm not sure which point release it'll be. Probably .2 or .3, that is, 6.2 or 6.3, because the first .1 release will solve the - it will remove SpinRite's use of the BIOS for AHCI controllers. So allow it to operate at the hardware level, but probably not, almost certainly not USB at that point release. But probably .2 or .3 it will then understand the USB hardware controllers, as well. So then it won't be necessary to access USB through the BIOS interface. Today it is.

So for that to work, because the BIOS doesn't understand plug and play, that is, you're not able to hot plug a USB as you will be once SpinRite is enumerating the USB bus on the fly, as it will be under .2 or .3, the secret is to have the phone powered up and connected to the computer when the hardware boots, not just the OS, but the BIOS itself because the BIOS comes up and looks at the USB ports and assigns them, technically it's INT 13 presences or appearances, which then SpinRite is able to see in order to run on the USB device. So it has to be a smartphone which is able to appear as a drive, but it also has to be present when you power on the PC in order to start SpinRite up. And if you achieve those two things, you can run SpinRite on a smartphone.

PADRE: There are a couple of other things that you can do to help along. Turn off all encryption. So if you've encrypted the storage, which you should, turn that off because I know at least on the OnePlus phones and the Samsung phones, when you do that, it requires an additional driver layer because that's what it's going to unencrypt. That's what allows your PC to act as a dumb USB device.

Steve: Ah, right.

PADRE: So turn that off. Secondly, it worked on my OnePlus One, but that was way back when my OnePlus One still had three versions of software ago. That allowed it to be seen as a plain stupid USB device. It was cheap storage. They took that out. So if you updated the software, they don't do that anymore. And I haven't tested trying to go back to a previous version of the software. I would assume it's probably not going to work. So right now it's kind of hit and miss. But as you mentioned, when you update SpinRite, that hopefully will help. Because it really did help. That OnePlus One was on its last leg. I couldn't get anything to work properly. It took forever to start applications. But after I did a Level 2, which is basically trash collection, it sped it up for about three or four months before it started slowing down again. And I just realized the memory on this is just worn out.

Steve: Is just getting worn out, yeah.

PADRE: Right, Steve. Let's close this loop. You know what, anytime we've had an episode together, we've never been able to do this because we always run really, really long. So I'm kind of happy we can finally do this.

Steve: So it was nice that we had a low security catastrophe week, relatively. So last week I turned our listeners on to the recently out of beta Quad 9 DNS service. It looks really good. Based on the available points of presence, I think it's 70 current points of presence of the DNS server using anycast, moving to - their plan is to have 180 by the end of 2018. For many people, it is very fast. But the main benefit is they are being very proactive about security. They return an NXDOMAIN error for any sketchy malware

domain that could hurt you.

And so the beauty is, if you configure your router to give all of the machines in your LAN 9.9.9.9 as its DNS server, you're prevented from having your browser look up the IP of a sketchy site; or, even if you went to a trusted site that had been compromised, that had some content it was trying to pull from the darkness of the Internet, again, that DNS lookup would be blocked. So results have been mixed in the first week of feedback, though largely positive.

So, for example, Jason Egan said: "With respect to Quad 9, do you suggest we utilize any secondary DNS, for example Google's 8.8.8.8, or use 9.9.9.9 solely?" He says: "I'm giving Quad 9 a go on my network now and was curious. Thanks."

And I would say yes. I heard one report of some sketchy service of Quad 9, but no details. As we know, nothing works without DNS, essentially, especially now that we're moving to HTTPS. You're getting a certificate whose domain name must match. So you can no longer substitute IPs for the domain name because the browser will say I'm trying to connect to this site by IP, which will cause a certificate mismatch; and browsers are really getting huffy about that, in order to protect their users.

So the problem is that, if there's ever a problem looking up Quad 9, systems are generally sticky. When the primary DNS fails, if the secondary one succeeds, they adaptively learn and will continue using the secondary, assuming that the primary is just not functional. So if you're relying on Quad 9 for security, giving it a nonsecure backup can cause your security to fail. On the other hand, maybe that's better than complete lack of access to the Internet.

I guess my advice, if you really want the security, would be not to provide a secondary and see how it goes. Be alert to the possibility that there might be a failure, and know what that means if it happens. But, I mean, this is a solid global network, and I only heard one report that I haven't tracked down of some sketchiness. And it may have been months ago. It may have been back before it was brought out of beta, and it's just no longer a problem. My advice, and it's what I have done, is just rely on Quad 9 exclusively and see how that goes.

PADRE: I was just wondering, I always put a backup. I'm typically - I'll use Google and OpenDNS. Those have been my favorite for years.

Steve: Right.

PADRE: So 8.8.8.8 and 208.67.222.222 or 220.220. And I know that there's a couple others that OpenDNS has.

Steve: Right.

PADRE: Can you think of - is there a way in OS X and in Windows, to determine which one it's using? Because I can't think of any command line that I have.

Steve: No. The only thing you could see would be if you were looking at the packets, you would initially see attempts to use the primary. When that didn't respond, then it sends out a bunch. Actually, when the primary fails, all of the other ones are simultaneously sent to, and then it sticky remembers the first one that responds, until you reboot. It doesn't make any record of it; but it just, in the software stack, it sets that then as its new preference. And subsequently you'll only see it asking the secondary or tertiary or

whatever DNS did respond.

PADRE: Yeah. And actually "PSchops" in the chatroom, he points out netstat might actually show you because it will show the active connections. So if you pinged something...

Steve: It's not a connection, though. And netstat won't show you UDP. It'll only show you TCP connections.

PADRE: So the way I could do it on this laptop right now is I would open up Wireshark, and I would just watch all my line traffic, send out a DNS query and just look at the address, see where it goes.

Steve: And that's why I think, until we have any reason to think that the service might be flaky, if you want the security, I'd just put in a primary DNS and see how that goes.

So Dan Kutka sent, he said: "Switched over to Quad 9 DNS and definitely see better performance than OpenDNS and SafeDNS I was using previously. A nice addition to my toolkit."

Kevin Sanders tweeted: "I'm taking it to the 9s. Even in rural Utah, Quad 9 is kicking it in gear." And he attached a screenshot of GRC's DNS Benchmark, which I'm sure our listeners know has pretty much taken over the industry. It took me quite a while to write it. Let's see, it's got a copyright of 2010, so seven years ago. But it is arguably the benchmark for DNS on the Internet. And this definitely shows that, I mean, it's amazing. He must be sitting on top of one of the Quad 9 points of presence.

PADRE: Yeah, that looks like an anomaly. I don't understand how that works.

Steve: Yeah, it's freaky because it's in number one with a much faster response. And then OpenDNS is in second place with one of the ones you mentioned, Padre, 208.67.220.222, then 8.8.4.4 and 8.8.8.8, both of the Googles are in third and fourth place. And then a whole bunch of OpenDNSes, and then several of the Level 3s. I think that was NTT, the 129.250.35.251 and .250. And then below it there was UltraDNS and some Level 3s. But if this, I mean, if what he gets is replicable, then that's a win.

But not everybody found that. Tim Grissom, he said: "In Orlando, 9.9.9.9. is slower by a factor of two than OpenDNS." And Chris Erice wrote: "Unfortunately, Quad 9 routes Seattle users to Palo Alto, California," he said, "verified by traceroute. Sticking with OpenDNS for now." So I wanted to just say that it certainly matters where you are. Apparently, if you're in Utah, next door to one of these points of presences, it's the fastest thing ever. And I heard from several people in Seattle who are listeners who all experienced the same thing that Chris did, which is that there wasn't a local point of presence for 9.9.9.9. in Seattle, so it didn't make sense yet. And this is why they're saying through 2018 they're going to be expanding to 280 GlobalPOPs, which hopefully will be more useful.

So I would just say use GRC's DNS Benchmark, see how it performs for you, add 9.9.9.9 to the list of DNS servers - you can do that by right-clicking in the context menu in the upper left of the benchmark, and you'll find a surprisingly long list of options which are there for power users that allow you to add servers, save an INI file which is then used by default when it restarts, and compare. See whether you are lucky enough to be close enough that it's at least at parity with other DNS options. And if not, you may have to wait a while. But you can check from time to time.

PADRE: Right. And that's why that benchmarking tool is fantastic because it does matter where you live and what provider you're connected to. Although I would guess that any of these, probably any of the top half of these results would be better than the response you're going to get from Comcast or Verizon because they've notoriously oversubscribed their DNS.

Steve: They're not even on the map. I mean, literally, they're not even on the list. You have to dig down in order to find them.

PADRE: Take a four-hopper over using Comcast DNS.

Steve: Right, right. Justin Alcorn responded by my confusion last week about what 9.9.9.10 was for because I commented that it was there. In my coverage of this I had picked up that it was unfiltered, that is, 9.9.9.9 would return NXDOMAIN. Anyway, Justin tweeted that .10 is for research. So you know if the NXDOMAIN returned by .9 is for blacklist reasons. So I guess that makes sense. That is, they're providing you with a differential query, one that is filtered and one that is not, so that you could, if you were a researcher, you could pull from both. And if you get a different response, then you know something definitive. And lastly...

PADRE: Now, I've got this map here, JammerB, if you can show this. This is old. This is way outdated, and this was before the IPv4 address exhaustion. So a lot of this is gone because they've given it back. But, yeah, there are whole domains here. There are whole ranges that were reserved in the early days of the Internet that are, quote/unquote, for "research." In fact, where's 45? Forty-five used to be all of Interop. We owned the entire thing. We had a Class A that was just given to us for our network. And it had so much space.

Steve: And 5, that was the Hamachi Class A that the Hamachi guy used because nobody else was using it.

PADRE: Yup, yup. And of course, if you're a good Internet citizen you've returned those so that they can be allocated, although they're now all gone. All gone gone.

Steve: And 9, does 9 show as IBM?

PADRE: Oh, yeah.

Steve: Does that show who has? Because 9 was always IBM's network. In fact, that's where we got 9.9.9.9 was they gave up that IP.

PADRE: Yeah. So you still had HP's old - DEC's old address range right here.

Steve: HP was 14 and 15; right?

PADRE: HP, yeah, 14, well, part of 14. All of 15. Ford had 19. MIT had 18. Apple had 17. Bell Labs had, what is this, looks like 12.

Steve: Wow.

PADRE: Their number is kind of weird on this graphic.

Steve: Wow.

PADRE: Yeah, just it's kind of fun. Look this up. This is the old xkcd comic that they made. Again, doesn't work anymore. Actually, IBM is 9. There's IBM. There's UK Mod, Cable TV - Cable TV? Hmm. And then the entire multicast range, which is fun. You know, I wonder, if I were to scan this again, who is where?

Steve: Yeah, it'd be fun to have an update.

PADRE: Yeah. Well, I mean, they can't make a map of IPv6 because that would be huge.

Steve: Yeah. So Nerds On Site are friends of the show. They were early sponsors back in the day. And I did want to note a tweet I saw from David Redekop, who said: "Great Quad 9 coverage. Just in case you're wondering, yes, of course we have DNSthingy support for it already." And DNSthingy is a cool solution that those guys up in Canada, the Nerds On Site folks, have created, DNSthingy.com. And sure enough, they've got a dropdown list that shows OpenDNS Family Shield, OpenDNS Home/Umbrella, and Quad 9. So using Quad 9 is as simple as selecting it from a dropdown list if you are a DNSthingy user. And DNSthingy offers all kinds of other benefits, as well, that we've talked about in the past.

David Hay said: "Hey, Steve. Have you yet had a chance to test Firefox Quantum and LastPass? If so, what's your verdict thus far?" So I've heard some controversy. I am a fan of Firefox Quantum on my Win7 machines. It is really fast. Lower memory consumption, and I like the way it looks, and I'm definitely using LastPass because I can't function without it. I'm not running across any problems.

I know that there was something about some incompatibilities. It may be some feature of Last Pass that I haven't encountered. But for what it's worth, it's working fine for me, and I'm not seeing any problems. And I imagine, if there was some startup glitch, that they will be resolving it. I do also remember saying or seeing that LastPass will be supporting the previous plugin technology for quite a while moving forward, even though they're moving now to the standard plugin paradigm that's being adopted cross browser, which I think is all good, all for the best. There will be overlap, so older versions should still be supported. But for what it's worth, I've not been having any trouble with it. And I can't operate without it.

PADRE: Steve, just a quick addition here. I just got a note from a TWiT Army, Simon Zerafa, who has learned that the Apple root account bug actually works remotely, as well.

Steve: [Gasp]

PADRE: So it works over VNC and Apple Remote Desktop. So, yeah, that suddenly went up in priority.

Steve: Boy. Boy. So you would - oh, my god. That would mean that, if you were exposing remote login, you can no longer rely on the machine's credentials to make that safe.

PADRE: Yup.

Steve: Boy.

PADRE: Yeah. So how about this? Until you get it patched, maybe turn off Remote

Desktop and disable VNC.

Steve: Boy, yeah. So not only physical local access, but even VNC and Remote Desktop. Boy.

PADRE: Ouch.

Steve: Yeah.

PADRE: Now I really want to go try that.

Steve: Yeah.

PADRE: I have an account for my sister's Mac. Hmm.

Steve: So Andrew, who tweets from @ISpaceCab, said: "@SGgrc On security, could you make a comparison between Telegram and Signal messaging apps?" Yes. Telegram is closed source, using a very bizarre, homegrown, untested, and unknown encryption protocol. Signal is open source, using well-tested industry-standard encryption, with completely documented and very clear additional features added. I would never use Telegram ever. Nobody who actually cares about security should. We've in the past covered well-known cryptographers saying, oh, my god, I mean, of Telegram. So I know people like it, its UI. It's candy. It looks nice. Fine. It's good enough because, I mean, if you're using a phone, you don't really have security anyway because phones can have other stuff in them that are intercepting keystrokes and messages and so forth.

So to some degree this is all an illusion, that is, messaging security. I've said to Leo, if you actually want security, you and who you want to have a conversation with need to strip naked, walk into the middle of Central Park, throw a black blanket over yourselves so that no one has any electronics, and no one can read your lips, and then whisper into each other's ears.

PADRE: Steve, that's a different kind of messaging that you're talking about.

Steve: Yes. That's where iMessage stands for "intimate messaging."

PADRE: Also I've heard that Telegram, the way that they encrypt it is they have a hard drive with built-in hardware encryption, and you put that into the program, and it works fine until you replace it with an unencrypted hard drive and [crosstalk] everything.

Steve: That's right. There's nothing worse than homegrown encryption. We cover it all the time. It's like, oh, Billy came up with this new cipher; and, oh, it scrambles the bits. Yeah, and it's super secure because it uses infinite bit cheese or some nonsense.

PADRE: We just can't show it to you because it's super secret.

Steve: Yes. And Telegram offered a large bounty, I think it was a million-dollar bounty for someone who cracks their encryption. No one has bothered because it doesn't matter. Signal, the guys at Signal know what they're doing, Marlinspike and company. I've covered it on the podcast. I did a couple podcasts about the Ratchet protocol and the cool things they did. It's the one you want to use, if security actually matters. If it doesn't, fine, use Telegram.

PADRE: That's a great endorsement there, Steve. Yeah, use Telegram if you don't care

about security.

Steve: If you don't care about - yeah, yeah.

PADRE: It's fine. It's perfect.

Steve: Yeah, yeah. I mean, it's good enough. It's going to scramble the bits. Someone is not going to be able to see what you're doing. I mean, no one has attacked its protocol because what they invented was just bizarro. They've just got arrows pointing around in circles around in the boxes. And it's like, okay, this really scrambles stuff up a lot. So it must be hard to unscramble. Okay, fine. But it's not based on any theory. There's no reason to trust it except it scrambles a lot. So, fine.

PADRE: So, Steve, this next one actually I am very interested in hearing your thoughts because the JTAG, the Intel JTAG vulnerability is pretty huge.

Steve: So this is Simon Zerafa, who tweeted again. He said: "Any thoughts on the Intel JTAG bug being deliberate? How useful would such a working attack be on any Intel platform to a three-letter agency?"

Now, I got a tweet from an Intel person who said that, in fact, JTAG over USB was a thing. That is, and I did some digging into this, like digging deeper. And the argument is that doing a separate JTAG interface is expensive because you need to have dedicated pins on the system on a chip on the SOC. You need to set things aside. You need connectors. You need traces. You need to open the box, blah blah blah.

So the idea is there's a rationale for JTAG over USB. I found a company that a couple years ago got a patent on some way of sharing a USB, an existing USB port which was still functional as USB, and simultaneously having it give JTAG access over USB so that you didn't need weird voltages or weird hardware. You were actually using the USB bus. I didn't dig into that any further. We're just going to have to wait, I think, until we learn at this forthcoming Black Hat, I think it's next month, it's in December, exactly what those guys at - I can't remember the name, it's like Precision Security or something like that. Those guys, they're going to tell us what they found.

Maybe Intel - the only thing that I can imagine is that Intel was being responsible. They recognized the security implications of exporting the JTAG interface, which is incredibly powerful, I mean, you can stop the processor. You can single-step it. You can suck the registers and memory out through this. I mean, it's unbelievably powerful. So they must have somehow implemented security like you have to send magic incantations using signed requests and date stamps and replay attack proof. And hopefully they did a really good job and just made a mistake. That is, there's just a bug, so that if you give it a packet that's too big, or you flip it upside down and send it in backwards, who knows. But they must have - hopefully they found an actual mistake, rather than JTAG just being exposed over USB.

PADRE: I'm actually afraid of that. I'm thinking they just figured no one is ever going to be able to figure this part out. I mean, who would be even looking for JTAG access over USB?

Steve: The ultimate obscurity reliance, which turns out to be insufficiently obscure.

PADRE: No, I'm with you. I'm hoping that they said, yeah, it was a very complicated process, and we needed to malform certain commands in order to get access for this.

Steve: Right, right.

PADRE: Yeah, I don't want it to be, yeah, they just assumed that no one would ever try to do this.

Steve: And we hooked up a JTAG debugger and, "Oh, look, we're in." Yikes.

PADRE: That would be a horrible presentation.

Steve: So Adam van Kuik said: "What was the 19-book series you were talking about earlier this year? You mentioned it a few times on Security Now!, and I believe you said you read it twice." For the record, it is called the Frontiers Saga. It is now - the plan is 75 books from the author, with whom I have established a dialogue, Ryk, R-Y-K, Brown. And it's at the top of my list of my absolute favorite science fiction of all time. And I've got a long list. Peter Hamilton is there. I created a PDF, and I'm going to have to update and put this first. The first arc is 15 books. He's now working on his second 15-book arc. He's six books into the second one, so what is that, 21 books total. And I am reading them a second time, and I'm now in the second book of the second arc. And soon I will be reading books I have never read before as I catch up to where he is for my reread. But I'm just loving it. So Frontiers Saga, the Frontiers Saga.

PADRE: Now, are they staying with the same group of characters over the 19 books?

Steve: Yes.

PADRE: So it's not just a universe, it's the actual same characters.

Steve: Yes. It is. It's future - and that's why this guy's so good. It's future history, a really interesting theme with good physics, interesting protagonists, people you really care about, really good characterization. He's created a fabulous world. And it's not the crazy world with people doing implants and really bizarro stuff. Just humans, humans set in the future who we care about. I just love it.

So also, and this comes to something that you were referring to earlier, Padre, Ove Karlsson sent: "Thank you, @SGgrc. Got bored on Friday, went scrolling through Netflix, and 'The Expanse' was on one of the cards." He said: "Remembering you raving about it on SN, so I gave it a go. Now a two-day and two-season binge later, I can't wait for Season 3." So "The Expanse," when I knew it was coming, I read the books, loved the books, and have loved it on Syfy. It is a topnotch production, which is very unusual for the Syfy channel, which is normally just, oh, my god, how many Sharknados must we tolerate in order to have an "Expanse"?

PADRE: You know what, if they give me - well, we are getting a Season 3. But if they give me a Season 4 of "The Expanse," they can make as many Sharknados as they want, and all is forgiven. We actually had the main author - there's co-authors - of "The Expanse" series, Daniel Abraham. He was here for "Triangulation."

Steve: Yay.

PADRE: I interviewed him. And one of the interesting things that comes out is Ty, the co-creator of - originally "The Expanse" was a board game. He was building a board game, and he had done all this research to make a board game.

Steve: Wow.

PADRE: And Daniel said, with all this research, you've got enough here to write a book. And so they started turning it into a book, an incredible book series. Now, here's the fun thing about "The Expanse," Steve. There's a lot of times where I have a beloved piece of work, especially in Syfy, that when they turn it into a series, it's okay. I mean, it's not great. It's not what I pictured. With this, they've cut so many things from the books, sort of combined different elements from the books, but it still feels right. And that's what you get when the authors are fully part of the project.

Steve: A classic one is "Under the Dome."

PADRE: Yes.

Steve: When that was coming out, I thought, oh, cool. So I read it - Stephen King. Read the book. Loved the book. The series was godawful. Oh, my lord.

PADRE: It was horrible. It was so bad.

Steve: Oh, it was just atrocious.

PADRE: No, but, I mean, if you read "The Expanse" series and then watch the series, they actually build on each other.

Steve: Yes, yes.

PADRE: And the casting was so perfect. Every major cast member, and even most of the minor parts, have been fantastically done. So kudos to them. If you're not watching "The Expanse," you're missing undoubtedly the best sci-fi currently on. I haven't been excited about a sci-fi project like this since the end of "Battlestar Galactica: The Reimagined" series.

Steve: I agree. Okay. This one's going to blow your mind, Padre. John Arundel said: "Mind. Blown. Pingfs is a filesystem that stores data in the Internet itself, as ICMP packets going to remote servers and back."

PADRE: This is great. Everything's in transit. It's always in transit.

Steve: Yes, yes, yes. So it's on GitHub, if anyone is interested. So one of the things we talked about years ago, back in our foundation of computing series, is the difficulty that early computer designers had of storage. I mean, when you think about it, how do you store a bit? Now, there's a thing called a flip-flop, which we've discussed, where basically you take two inverters, and you connect the output of one to the input of the other, and the output of the other to the input of the one. And that's stable. So that, if the input of the first is high, its output is low, which goes into the second inverter. So its output is high, which folds around and keeps the input of the first one high. So it's stable. If you force it into its other configuration, that is, invert the lines, then that's stable, too. So that's two inverters cross-coupled, as it's called. And you could do that with a tube, and they did.

So they had tubes. A pair of tubes could store one bit or a relay. You could have a relay that would latch itself, or two relays that were inverters that would also be a flip-flop. But those are large. They consume a lot of energy. And they store one bit. So storing bits was a huge problem early on. One of the solutions that was actually developed was a mercury delay line, where it was a large tube - technically a column, but it's on its side -

a tube of mercury with an ultrasonic transducer on one side and a microphone on the other. And the speed of sound through the mercury was used to hold data. That is, this was like a recirculating acoustic delay line, and the output was fed back into the input. So there were enough bits in there to be practical.

And so that's what I was put in mind of when I read about this Ping File System. You know, a ping isn't just an event. It actually can be a full-size UDP packet or IP packet. So it can store, what, 1,500 bytes, or is it bits? 1,500 bits.

PADRE: It's mostly empty, but it doesn't have to be empty.

Steve: Doesn't have to be. Typically - exactly. Typically it's just nothing. But you could store bits. And you send it off, and the goal of the recipient is to return it to you. So you send it as a pin query or an echo request, and it sends you an echo reply. And so you could send data off that you then forget, and you rely on the remote end sending it back to remind you of what that data is. Essentially, if you were crazy, you could build a file system. And you would need redundancy and error correction because packets can be lost.

PADRE: Right, right. You can lose a packet. You'd have to send multiple copies of the same packet in different directions.

Steve: Off in different directions. And you keep them going back out again. And then they'd be recirculating, and so then if you wanted to look for some data, you'd have to wait for the packet with that address to come back to you, get the data out of it, send it back out again. Maybe you'd do a read-modify-write cycle, which is what we had to do in the core memory days. So when that comes in, you would get the data, see if you wanted to change it; if so, change it and then send it back out again. And in theory you could do a file system on the Internet. Anyway, I knew you'd get a kick out of that.

PADRE: For those geeks out there who think that this sounds familiar, that's because you've probably remembered this from I think it was Season 6, Episode 4, the 130th episode of Star Trek: The Next Generation called "Relics," when Scotty was discovered crashed on the outside of a Dyson sphere.

Steve: Yes.

PADRE: He and his co-engineer had holed themselves up...

Steve: In the pattern buffer.

PADRE: ...in the transporter buffer, right. But because the pattern buffer will eventually lose its pattern integrity, they did a cycle between the dematerialization and the rematerialization cycles back and forth, back and forth. So I'm not sure why that came into my brain.

Steve: That's good. That's good.

PADRE: But you're right. Mind blown. Thank you.

Steve: So JMWhitty said: "When you talk about Apple versus FBI regarding golden keys, you often do not mention the impact to the non-Apples of the world." And I wanted to say yes, that's true because among the things Apple has done is to provide, typically annually, these amazing security whitepapers where, while they don't go into code level,

they really give a very useful overview of the architecture of their security, enough that we've been able to do podcasts on it several times. The other solutions are relatively closed.

I talk about Apple because I know enough about the architecture of their security to put the impact of encryption in context that Apple has provided. But we just don't have - the other manufacturers are just black boxes. It's like, oh, yeah, we have security. Trust us. And so if there were other manufacturers issuing the same level of detail, I'd love to know, first of all, to know more about it, to share it with our listeners, and then be able to create some context. But we just don't have the information.

Mike Synan says: "Any suggestion for network firewall hardware not containing IME? Or how to build an uncompromised network firewall without the hardware hack discussed two episodes ago?" And of course the answer is a Raspberry Pi, is to do non-Intel. You want to stay away from Intel. And you can certainly do a Raspberry Pi on a non-Intel platform. So Padre, your suggestion?

PADRE: Well, I would go with the - so SonicWALL uses the Cavium processor, which again is - it's not Intel, doesn't have any Intel architecture. It's ridiculously fast. It's used in all of the SonicWALL SOHO, all the way up to their enterprise class. They just add more processors to get more speed. They just entered into an agreement with Cray, that's right. Cray is going to be using Cavium processors in the new supercomputer blades.

Steve: My lord.

PADRE: Yeah. So they've done a little bit of development over the years. You can find lower end Cavium processors in bargain bins now, from the old Dell hardware. And it's relatively easy to develop. You can actually...

Steve: Are they ARM architecture?

PADRE: It's an ARM offshoot?

Steve: So probably an ARM license.

PADRE: Right.

Steve: And then they did their own thing. Wow.

PADRE: Right. But it's ridiculously fast. Their silicon is fantastic. So if you wanted to make a high performance one, that would be my suggestion.

Steve: Cool, thank you. So Kyle from Craig Consulting said: "Hey, Steve. If you think watching disk defragmentation is mesmerizing, you should try watching a 3D printer at work. Very soothing."

PADRE: Yes.

Steve: I think it was last week or the week before, our Picture of the Week had the caption "I Defragged My Zebra." And you know how a zebra is all covered with black and white. This showed the front half black and the second, the hind end, all white because of course it pushed all of the color to either side. And so we had fun with that. Anyway, so Kyle is suggesting 3D printing. And, yes, and I'm sure you know, Padre, it is certainly

fascinating to watch a 3D printer at work.

PADRE: This is actually the 3D printer I just tested at home. This is the Monoprice Maker Select Plus. And it is, especially if you watch it in time lapse. I love how this looks because you see it build up layer after layer after layer. I mean, what kind of a geek wouldn't love just to see that all the time?

Steve: Yup. Very cool. And finally, Will Springer said: "On Security Now! you and Leo were discussing the trend and importance of pushing updates to IoT devices to patch firmware vulnerabilities. Would that functionality create a new vector for a party to push malicious firmware to these devices? Thank you." So it creates the issue, but we have the solution. And coincidentally, we were just discussing a failure in implementation in those 54 HP enterprise-class printers. It is trivial to do this correctly, and that is, any incoming firmware has its signature carefully tested before it's deployed.

For example, I do this with SQRL. The SQRL client has a mature, finally mature, automatic update capability. When it is informed that there's a new version, it downloads the version to a staging area, and it performs an Authenticode signature verification on the newly downloaded code, before that code ever has a chance to run. By definition, you cannot have software perform an Authenticode test on itself because it could be subverted, in which case it would say, "Oh, I'm fine."

So you have to have a third party look over at the code you're considering deploying and verify it before it ever runs. The SQRL client does that; an IoT device could do that. And there's nothing reverse-engineerable about it if it's done right because the IoT devices would only contain the public key which could be used to verify the signature that could only be created by the matching private key that, assuming that things are done correctly, never leave the control of the entity signing the firmware updates.

PADRE: Right.

Steve: So everything, all the technology, all the infrastructure, all the capability is in place for doing this securely. It's got to be implemented correctly. HP didn't. But a light bulb might be able to.

PADRE: Yeah. What we find is, as you mentioned, we have the technology. We know how to make this process secure. But you have to start using it in the design stage.

Steve: Yes, early, early.

PADRE: The problem is security is often added after everything's been created. And then you kind of fudge the security in order to make it work properly. Well, if you know that you need a couple of things - you need a signature, a bulletproof signature. You also need a way to change the signature in case it gets compromised, so that's something that a lot of these manufacturers don't take into account. If they lose the magic key, what's going to happen? So do you have a way to fix that and repair damage that's been done? And then you just have to make sure that there are no shortcuts to that. It's relatively simple when you put it that way. It's just making a manufacturer actually stay to security best practices that tends to be the problem.

Steve: Yeah.

PADRE: Well, you know what, Steve, that's actually a positive note. I don't know the last time I've ended a Security Now! feeling better about security.

Steve: Yay. Yes, there is hope. There is a future.

PADRE: There is always hope.

Steve: And we can get there, yes.

PADRE: We can get there. Folks, of course Steve Gibson is the genius behind Gibson Research. He provides us such products as ShieldsUP!, of course SpinRite, and soon coming SQRL. I know that you've toned down the SQRL talk over the last couple weeks, but I'm still very excited to see what you come up with.

Steve: We are getting very close. It's funny you mentioned ShieldsUP! - 99,993,060 shields tested. So we are, what, less than 7,000 tests away, so that's about two days at our current testing rate from crossing the hundred million mark.

PADRE: You could have one of those McDonald's signs. I can remember when I was kid, I still remember...

Steve: Hundred million served, yup.

PADRE: I saw, like, 10 million served.

Steve: Yup.

PADRE: And then it was 100 million served, and a billion. Now they don't even do it anymore.

Steve: No.

PADRE: But Steve, can you tell the folks at home where they can find you? I mean, of course they need to go to Gibson Research. But where are all of the places they can see the wonderful work that you do?

Steve: GRC.com. That's where everything lives. We've got a menu at the top, and you can browse around there to your heart's content. And I can't predict when SQRL will be ready. It's being tested now by a small group in the SQRL newsgroup. I'm going to stage it to the entire community that hangs out at the new server, which is a much bigger group, before - we have about, I don't know, we have hundreds of thousands of listeners to this podcast, and I want to make sure it's nailed down before it goes that big. But I'm working on it full-time, so I'm very excited about it.

And of course SpinRite is, as Leo always reminds everybody, is what pays my bills and gives me the freedom to do this. So I certainly appreciate everybody helping themselves with SpinRite. And as soon as I get SQRL off, I'm back to working on the first point release of SpinRite 6, which I can't wait to get back to.

PADRE: Someone once told me that you're the George R. R. Martin of the security world, where people are just hoping that you sit in a room somewhere and just keep writing, just write and write and write, so we can have what we want.

Steve: I love to do it.

PADRE: Steve, it's been an absolute pleasure. I know this was a bit of a surprise. I guess

no one told you that we were going to be together.

Steve: Worked out great.

PADRE: Anytime Leo goes on vacation, I am so happy to jump in the booth because this is one of my favorite sets of hours that I do here on the TWiT TV Network. It's been an absolute pleasure to be here and talk about all the important things that our digital lives require.

Steve: Great to be with you, Padre. And until next time.

PADRE: That does it for this episode of Security Now!. Don't forget that we're live here on the TWiT TV Network every Tuesday, 13:30 Pacific time. Now, Steve will be here to inject you with some healthy paranoia to help you understand the wonderful world of security. That's for those of us who don't have his genius IQ.

You can find all of our shows here at TWiT.tv/sn, as well as iTunes, Stitcher, and wherever fine podcasts are aggregated. You can also find high-quality audio downloads at GRC.com, which is also where you'll find everything that GRC offers, including SpinRite; ShieldsUP!; and, coming soon, SQRL. He's sort of the Szechuan sauce of the security world. I'm Father Robert Ballecer, the Digital Jesuit, saying that if you want to keep your data into the future, you need to start thinking Security Now!.