Listener Feedback #191

Description: Father Robert Ballecer and I discuss the week's major security events and discuss questions and comments from listeners of previous episodes. We tie up loose ends, explore a wide range of topics that are too small to fill their own episode, clarify any confusion from previous installments, and present real world 'application notes' for any of the security technologies and issues we have previously discussed.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-463.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-463-lq.mp3

SHOW TEASE: It's time for Security Now! with Steve Gibson. Microsoft clamps down on XP hackers, then saves the Internet by breaking it. Also a pizza-making robot; Oracle gets cagey on Java; your WiFi light bulbs are leaking security; and your questions, Steve's answers. Security Now! is next.

FATHER ROBERT BALLECER: This is Security Now!, Episode 463, recorded Tuesday, July 8th, 2014: Your questions, Steve's answers, #191.

It's time for Security Now!, the show that covers your privacy and security online. I'm Father Robert Ballecer, the Digital Jesuit, in for Leo Laporte. And of course guiding us through the swampy mess that is electronic privacy is the one, the only, the purveyor of proper packeting, Mr. Steve Gibson. Steven, so good to see you. I had so much fun last week that when they told me that I could do it one more week I was just - I was beside myself.

Steve Gibson: I was quite pleased, too. I think it's great. And I was thinking, you know, you are technically a Digitized Jesuit because all we're seeing is the result of that digitization process.

FR. ROBERT: Absolutely.

Steve: So digital and digitized.

FR. ROBERT: And people have asked, if I'm the Digital Jesuit, where's the Analog Jesuit? There actually is a Jesuit who goes by Analog Jesuit. I think he started his moniker after me, so I'm going to sue him.

Steve: Okay. Now, you wanted to talk about pizza.

FR. ROBERT: I did. I did. Okay. This is silly. I think the only way I could possibly tie this

into security is when you get into a late-night session of packet prioritization and scanning, what kind of food do you look for? I mean, for me it's always pizza. If you could jump over to my machine, this is a machine that some family filmed in Italy. It's a pizza vending machine. And when I mean pizza vending machine, it doesn't microwave a frozen pizza. When you choose your selection, within two and a half minutes it will make, from fresh ingredients, a pizza, and then cook it for you.

Steve: Wow.

FR. ROBERT: And you get to watch it. I mean, this is like food theater. I think that's basically what this thing is. So the reason why I wanted to show this is I want to enlist your help, if you can maybe convince Leo...

Steve: Oh, and so we're actually looking through the front pane of this machine. So they've deliberately made it visible so that people can see, like, oh, like there's the flour for the dough at the top.

FR. ROBERT: Absolutely.

Steve: And so it goes down and kneads the dough and flattens it out and builds your pizza. Wow.

FR. ROBERT: And isn't that part of the fun of going to like an authentic pizza parlor? You get to see the pizza being made? I mean, it's one thing to have it delivered. But if you can see all the different processes that go into turning raw dough and water into a pizza, well, that's just kind of cool.

Steve: This thing could absolutely survive in Silicon Valley. It would be in constant use.

FR. ROBERT: Now, one thing, I can't find the clip, but evidently, because it does have a menu screen, it has a touch screen that you get to choose your pizza and your ingredients, it is running Windows XP Embedded. So there's another security angle. I want Leo to get one so we can attempt to hammer on it, maybe break in and hack ourselves some free pizza.

Steve: Oh, you guys definitely need one. Now, if it's a legitimate XP Embedded, it's secure through 2019. So you don't have to worry about hackers messing up your ingredients, like giving you too much pepperoni. Not that that would ever really be a problem.

FR. ROBERT: Yeah, I don't think there is such a thing as too much. Too much pepperoni is like too much bacon.

Steve: I've never had too much pepperoni, yes.

FR. ROBERT: All right, now let's get onto something that's a bit more actual security, in the real world security. And that is a report from Incapsula. Now, Incapsula Networks serves about 20,000 websites, and they've been collecting stats, I think since 2010 is when they started looking at the amount of traffic pulling through their networks, specifically paying attention to the percentage of bots. Now, according to their latest study, which was based on 1.45 billion visits over a 90-day period, which geographically represented the entire slate of the world's countries, 249 countries, they found that 61+% of all website traffic is bots. That kind of blew me away. I know you've covered this on Security Now! previously. But I think it's a good reminder that most of the traffic

flowing on the Internet is not from human interaction.

Steve: It's just, well, and when you think about all of the things we're doing now which are services that are becoming uncoupled from, like, web browsing, there's so much more going on. And of course search engines are, I mean, we sort of take it for granted that we can put any phrase in that we can think of, basically ask the Internet a question, and we just get relevant results. And I remember, and I'm sure you do, too, a time before Google. And actually the search engine I think we were using before then was AltaVista. That was my favorite pre-Google search engine.

FR. ROBERT: I loved AltaVista.

Steve: Yeah. But the idea was there was all this stuff there, but you couldn't find it. Now we just take for granted the fact that, if it's there, we can find it. But something has to have gone there first to browse around and pull all that information together. And of course those are bots. Those are spiders that are out link-following the entire Internet. And I think, I mean, there's almost never a time when Google's not crawling GRC. I don't normally keep logs. I turn logs on if there's a problem. But just sort of being pro-privacy, if I don't need them, I don't have them. And of course I delete them when I'm through.

But during times when something weird's been going on, and I've turned logging on, I see Google bots walking around inside GRC. And I know when I bring a page up, I'm frequently surprised how quickly it appears in Google's index, which tells me since I put the page online Google found it, meaning that somehow something came to my site and saw the page and added it to the index. And it's a matter of hours. Which means, I mean, and my site's not super, super popular because Google does change bot visit frequency based on the popularity of the site. So they're, like, constantly crawling news sites whose pages are constantly changing. Because that's the other thing that Google looks at, and I'm sure other search engines do, too, is when they revisit a page, is it the same. And so that sort of lessens their urgency on checking in more frequently. So there's complex algorithms.

But basically we can think in terms of the Internet is constantly under the traffic of bots. And as you said, we first mentioned this back in mid-December, and I remember using the phrase "The bots are winning" because there are more bots than there are people now. I mean, like, by far. Well, and if nothing else, they're tireless. They're out there constantly scouring. And didn't that chart also - or you had some numbers about malicious bots versus just happy spidering bots.

FR. ROBERT: Absolutely. Now, so Spooky Spy in the chatroom is saying "Don't fear the spider." And, yeah, absolutely, don't fear what you're talking about, those search engine bots that are going out there, doing us a service. They're making things findable.

Steve: I want them.

FR. ROBERT: Yeah, exactly.

Steve: They're allowing people to find stuff on my site, yeah.

FR. ROBERT: Right, right. And according to this survey, most of those bots are actually good bots, quote/unquote "good bots." So they are, the search engine bots. And then you've also got a percentage of those bots that are kind of neutral bots. They're the SEO bots. So companies that are offering services for SEO, they've got a way of poking around to find out what they should be labeling their content with. It's kind of the gray

area. It's not good; it's not bad. It might throw off some results. They've also found that there's actually been a decrease in the percentage of traffic related to spam. From 2012 to 2013, they dropped about 1.5%, from 2% to 0.5%. So that's a good thing. However...

Steve: Wow, yeah.

FR. ROBERT: ...they've also seen a rise in the percentage of malicious bots. About 31% of that traffic is directly from malicious bots. And within that 31% you've got your regular DDoS bots. You've got your bots that are aimed towards NTP and DNS amplification attacks. And then you've got 8% of it that they're calling "other," which I interpreted as that's the advanced persistent threat bots, the ones that we know they're there; we don't know what they're doing. They're not exhibiting any behavior that will trip off any signatures.

But it does seem that they have some sort of - they're allowing some sort of attack vector. And that's the part that scares me. Not that there's a lot of bots. There's always going to be a lot of bots. Not that there's attack bots. Of course there are attack bots. But there are bots that security professionals are looking at and saying, we don't know yet what this is for. And we probably won't know until they do their thing.

Steve: Just listening to you and imagining turning the clock back to the beginning of the Internet, it sounds like science fiction. I mean, we're talking about this network that was originally built for people to click links to look at web pages to deliver information. Look how far it's gone. And it's gone into the realm of science fiction, of there being entities inhabiting - I mean, "inhabit" is not the wrong word to use - inhabiting this network by installing themselves on autonomous machines so that they can then reproduce and find other machines, I mean, they're inhabiting the Internet. And we're sure they're not self-aware. Yet, as you said, they have purpose, which unless they are captured and dissected, we can't necessarily divine just from looking at their behavior. So, I mean, it's science fiction.

FR. ROBERT: You're kind of freaking me out, Steve. What if these are bots that were created by other bots, like second-generation bots? Now we're getting into science fiction.

Steve: Well, technically there are polymorphic bots, where one bot will encrypt itself so that it cannot be recognized, and then send that encrypted version out. So bots are reproducing deliberate variations of themselves that will not be identifiable where they might be. So, I mean, there is reproductive behavior and, like, genetic variation to see whether it survives better in the wild. And of course the ones that don't get caught tend to perpetuate their pattern. So it's evolution on the Internet.

FR. ROBERT: Again, that's what's got me worried, the ones that don't get caught. Because we had Raphael Mudge on This Week in Enterprise Tech a few weeks ago. And he was showing us some of the research that he's been doing into advanced persistent threats. And he was showing us bots and malware that could perfectly mimic regular network traffic so that, even if you were watching it, you wouldn't be able to figure out that it was doing something bad.

So, and this study would not be able to see any of those. It would not include any of that traffic because that just looks like DNS traffic or FTP traffic. And it disguises itself so that it sends it out with the rest of the traffic, so you're not getting a traffic spike, so you can't find it that way. And the amount of intelligence that is being put into creating some of these bots, it defies imagination. I mean, they've really got it down to a science where, if you run any sort of large network, you really don't know if you've got

something running inside.

Steve: I was just looking this morning, I was doing a packet capture of some traffic from Paul Thurrott's site, the WinSuperSite for Windows, following up on one of the questions in today's Q&A, because it was about whether Paul's forms were secure or not relative to LastPass. And two things. One was the cookie content was complete gibberish. I mean, here I am, a security-conscious, security-aware person, inspecting the traffic, and I expect the cookie content to be gibberish. I recognize that they're supposed to be opaque tokens, and I'm a little annoyed by how much there is and how big it's grown. But it's not supposed to be understandable. So it passes my scrutiny. I go, oh, yeah, look at that disgustingly huge cookie blob. But I don't look any further.

And similarly, remember now that the web is running on JavaScript more and more. And JavaScript is code. And so when huge wads of it pass by, you just sort of go, okay, well, I mean, who's going to take the time to deobfuscate it and work through exactly what it does? And when you do, it probably has links to other JavaScript pieces that it pulls together from other, I mean, so it almost becomes an insurmountable task if you absolutely have to understand everything about what something is doing because, just sort of through convenience, people are pulling pieces together.

There are sites where, when I enable scripting, suddenly NoScript will show 30 other domains which it's now blocking which were revealed. But when I allowed the main script to run, that was invoking all kinds of other stuff, I mean, analytics and tracking and little blobs coming in, like online JavaScript pieces of toolkits that the web developers have just used for convenience. And they're not malicious. But they're just, like, pouring in from every direction. And so it's like, oh, my god, how can anyone actually know what is going on? And the fact is, it really has gotten away from us. At this point, it's not clearly doing overt damage. But are we in control any longer? No. We've lost control of this thing that we've built.

FR. ROBERT: I think that's enough freaking out of the audience. They're starting to hate us in the chatroom. All I can say is, everyone, after the show, go watch "The Matrix," the first one, not the next two, and just realize that that's probably next year.

Steve: So we've got a Patch Tuesday to talk about. Microsoft was also - they also fumbled a web domain takedown. I just didn't get it into our show last week, so I wanted to cover it because it was just an interesting point. Some news about Oracle's maybe ending Java support under XP, but maybe not. Google just finding some more unauthorized Google certificates in the wild. A cautionary tale, I mean, we've already set people up for some caution, but another one about just sort of the general problem I think we're going to have with the so-called "Internet of Things." A little bit of follow-up on last week's announcement of our intentions to open up the whole topic of cloud storage solutions. And a bunch of miscellaneous stuff, but not too much. And then a Q&A. So a great podcast today for #463.

FR. ROBERT: That's an absolutely great lineup. That's Mr. Steve Gibson. I'm sorry, what's your Twitter address again? I always want to put Steve Gibson, but I know it's @SGgrc.

Steve: Yup.

FR. ROBERT: That's right, the man who provides us all the tools that we need to keep everything running, including ShieldsUP! and my personal favorite, SpinRite. Now, Steve, I talked about this last week, that I was having some issues with an SSD. And I still, I want you to explain to me how this actually works. Someone told me that the throughput issue I was having with some of my older Samsung SSDs could be solved by running

SpinRite in Level 2. And I was thinking in the back of my head, I'm like, this is ridiculous, it's an SSD, it's not a rotating drive, that's stupid. I ran it, and by golly, it worked. How does this work? What magic did you bake into SpinRite to revive SSDs?

Steve: What we would like to believe is that, because they're solid-state, they're like RAM, where when you write the data, you are sure you're going to get it back. Now, even RAM is not perfect, which is why there's parity in one case, just to find an error when it's read back. We have parity and non-parity RAM. And there's also ECC on RAM. When really, for example, space shuttle RAM will have ECC because you can have, at the quantum level, you can have bits that don't read back correctly.

Well, the situation is far worse with SSD. Essentially, the technology of SSD is like dynamic RAM, where in order to get the bits small enough, dynamic RAM uses capacitors to store the bits as ones and zeroes, just uses electrostatic charge. And that tends to bleed off over time just due to leakage because the cells are so small, the capacitance isn't large enough. And it's a series of tradeoffs. And that's why you have to do so-called "refreshing" of dynamic RAM. You've got to come back and read it before it's bled out, before the data has sort of leaked away to a point where you can no longer differentiate the ones and the zeroes. So you keep coming back and reading it and rewriting it to, like, to recharge the little capacitors with their data.

Now, freaky as it sounds, that's the same technology as in SSDs. It's a much slower leakage, but it's still doing that. And so it actually - an SSD is just a huge plantation of little capacitors where charge is essentially stranded out on a plateau, and a field effect transistor is able to sense the field, the electrostatic field created by that charge. But over time, these cells weaken. And the important thing to understand is that, if engineers only needed to make a 1K SSD, oh, my god, it would be bulletproof, absolutely reliable. We could do that because, with so few bits, the bits could be so large that they could be reliable.

But the world runs on economics and competition. And so we've got multiple vendors who are competing with each other to get the highest density and the lowest cost. What that means is the smallest bits. So just as we have pushed hard disks to the point where they are now using error correction all the time, that is, you can't, you often cannot read a sector correctly on a hard disk because the engineers said, well, you really don't have to. We can correct it, as long as it's not too bad. So although, I mean, it's really cringeworthy because this is our data and we care about it, but the same thing has happened with SSDs. The engineers have pushed them so far that they're operating more in an analog fashion, not just one and zero, but somewhere between one and zero.

So what SpinRite is able to do is it's able to turn off, by talking to the drive, it turns off some of the sort of, oh, don't worry about this, we'll take care of it stuff, in order to show the drive when it has a problem which it would otherwise ignore. That forces the drive to address the fact that this area is no longer stored safely, that is, some of the bits are beginning to wander toward an indeterminate state. And that causes the drive to rewrite them firmly, but that takes a little bit of time.

See, SSDs don't write very quickly, as we know. They only read quickly. The reason is they actually have - there's a layer of insulation between this little floating island, and they have to ramp up a high voltage and push electrons through an insulator using high voltage in order to recharge that island. Well, that's why there's a limited number of writes that an SSD can do because every time you push electrons through that insulation, it weakens the structure of the insulation a little bit. And so that creates a lifetime on the number of times you can do that.

But so essentially what SpinRite does is it allows the SSD to be more picky. And instead of being lazy and using error correction to fix the sector which is becoming weak and taking more time, it says, no, let's, like right now, let's fix this. And so the SSD rewrites that sector which was using error correction so that it no longer needs it, which then speeds up the execution in the future. So, I mean, there is - we just plug these things in and format them and go. But there's an incredible amount of technology under the covers.

FR. ROBERT: Now that you say it, it makes so much sense. Rather than the drive trying to fix the errors on the fly as it's reading and writing, SpinRite just goes in and says, you know what, I'm going to fix everything for you, and you're golden. And that would explain why suddenly I get all my performance back. It wasn't necessarily that the drive was damaged, it's just that the drive was busy. But because the way that SSD drives work, it didn't want to do all that maintenance because in doing that, it's actually wearing itself out. I like that. That's fantastic.

Steve: Right. It just wanted to defer that.

FR. ROBERT: Now, Steve, I would never have figured that out. So people need to go and get SpinRite. This actually is brilliant because I had thought that, once the era of the SSD had been upon us, that SpinRite, which has been a great friend to me in my troubleshooting days, would go the way of the floppy disk. But this just shows you that you never know when the technology from a brilliant man might come in handy. Where can they find SpinRite, if they want to get a copy for their SSD?

Steve: GRC.com/SpinRite, or you can put SpinRite, S-P-I-N-R-I-T-E, into Google. It'll take you there. And it's been now more than 20 years we've been selling it. I am working on wrapping up the work on SQRL, and I have some fun news about that. I got a tweet from someone this morning who's been playing with it because people's implementations are beginning to come alive. So that gets done. Then I get back to SpinRite. I'll be producing 6.1, which will be a free update for everyone who has SpinRite 6. We are at that point going to stop bringing everybody else along from 10 years ago or 20 years ago, really, because we've had SpinRite 6 now for 10 years, and that seems like enough time for people to update, and it simplifies things for us if we're not continually trying to bring people forward from 20 years ago.

FR. ROBERT: No, Steve. No, I bought my copy 20-plus years ago. It should work today. We've got no analog for that, thankfully.

Steve: Exactly. So anyway, it's because, as you say, there's a clear future for it that I fully intend to then move to SpinRite 7. There will be something following 6. But I want to get SpinRite 6 caught up to date with the latest BIOSes, compatibility with Mac and UEFI stuff. And also so that people don't have to, like, change BIOS settings in order to get it to go. It'll just run in whatever your environment is, so it'll just be easier to use. I want to do that sort of as an interim measure before I start in on 7 because I'm intending now to scrap the codebase of 6 and start from scratch on SpinRite 7. So that's not going to be an update. That's going to be a complete...

FR. ROBERT: Clean sweep.

Steve: ...from scratch rewrite because it's been 20 years. I think it's time to start over.

FR. ROBERT: Well, you've learned a few things in the last 20 years, I'd say, things that could probably be put to good use inside SpinRite.

Steve: Well, and SpinRite still runs on the text page. I mean, it's a text UI.

FR. ROBERT: There's nothing wrong with a text page.

Steve: It works just fine. But I think it's time to give people buttons and dropdown menus and things that they're used to.

FR. ROBERT: I have it running on a screen in the back of my lab. It just reminds me. I see that little progress bar, it's like hope, hope flowing on my computer. Now, Steve, we've got people clamoring in the chatroom to talk a little bit about Patch Tuesday. You want to start us off?

Steve: Well, yeah. So as I mentioned last week, since last Tuesday was July 1st, this is the earliest Second Tuesday of the month possible. And I haven't seen the patches yet. I got the email from Microsoft after they announced, as we talked about last week, that they were no longer going to be sending email because of the Canadian antispam legislation. I mean, it sounded like that was just some strange misfire somehow at their end because they quickly reversed themselves. I got first an email saying, oh, by the way, we're going to resume sending you emails because we changed our mind. And then I got the announcement a few hours ago about the content of this Patch Tuesday.

There were a ton of fixes for Internet Explorer. Essentially it boils down to two critical lumps of patches, three that Microsoft rates as important, and one moderate. Of the two critical ones, one was a remote code execution vulnerability in IE. Microsoft wrote: "This security update resolves one publicly disclosed vulnerability and 23 privately reported vulnerabilities in Internet Explorer." So a total of 24 in that one bundle. Of course they called it their "cumulative security update for IE," as they always do.

And as usual, Microsoft says: "The most severe of these vulnerabilities could allow remote code execution if a user views a specially crafted web page using Internet Explorer. An attacker who successfully exploited these vulnerabilities could gain the same user rights as the current user." So as always, you're much better off using your PC not with admin privileges, as a standard user, although it's more hassle. Normally you only need to have admin privileges, of course, to install drivers and to install some software. So, boy, over and over and over we see examples where not running as an admin user saves you. It's just the software that can do bad things can only do them on your behalf. And so if you deliberately restricted what you yourself can do, you're curtailing what malicious software can do.

FR. ROBERT: Now, Steve, that 037, that's just an extension of the browser pivot that we saw, like, six months ago; right? Is that the same flaw, or is this a new wrinkle on that flaw?

Steve: I've not looked any further than just - about half of these 23 looked like they were remote code executions. And they said they were a memory corruption vulnerability. So I just didn't go any further. It's like, okay, fine. I mean, everyone who listens to this podcast knows that IE should really not be your first browser of choice anyway. You should be using Firefox or Chrome and preferably running with scripts disabled and selectively enabling scripting where you know you need it, just for the sake of getting a page to run that doesn't run without scripting.

So then 038 is the second one, the second of the critical ones, which is also a remote code execution vulnerability. Again, we don't know much about this. All we're told at this point is that a Windows Journal file, a maliciously crafted Windows Journal file can be

used to perpetrate a remote code execution. So again, we don't know anything more about it than that's all they're saying, that it would allow a bad guy to somehow get privileges again of the current user if you open a specially crafted Journal file.

And then the three important things, there's a vulnerability in On-Screen Keyboard that could allow a privilege elevation; vulnerability in Ancillary Function Driver, AFD, which is one of the DLLs I've seen in Windows for years, that could allow elevation of privilege; and a DirectShow problem, as well. So again, it's time to update, everyone who still has Windows which you can update.

FR. ROBERT: Yeah, if you're in XP it's time to just cry, just cry a little bit.

Steve: Yeah, and I have some - I'll follow up on a question that we're going to cover today about that because it looks to me like Microsoft has already disabled the little trick that has been used since April.

FR. ROBERT: I did that. I set one of my XP laptops to report that it was Windows Embedded. And it worked. It was amazing. So you're telling me I can't do that anymore.

Steve: I think that jig is up. I think we got three, what, April, May, and June. We got three months' worth of extension. It would have been nice if it had been five years, but no.

So, okay. Microsoft did something that caught the attention of the press and the industry only because it was really needlessly heavy-handed. And it feels to me like there's some "uncoordination" among Microsoft. We saw this, for example, with this weird email announcement where they were going to stop sending the security announcements. And then within a couple days they said, oh, never mind, we're not going to after all.

Similarly, they essentially got a court order to take the 22 domains from a dynamic DNS provider without any notice. All of their four million website domains went dark. So this also, this I was mentioning we didn't get to talk about last week, so I wanted to. It's NoIP.com is the site. And so at the beginning of the week of the Fourth of July week, Monday, June 30th, they announced: We want to update all of our loyal customers about the service outages that many of you are experiencing today." And frankly, it's all of you.

They wrote: "It's not a technical issue. This morning Microsoft served a federal court order and seized 22 of our most commonly used domains because they claim that some of the subdomains underneath the primary domains had been abused by creators of malware. We were very surprised by this. We have a long history of proactively working with other companies when cases of alleged malicious activity have been reported to us. Unfortunately, Microsoft never contacted us or asked us to block any subdomains, even though we have an open line of communication with Microsoft's corporate executives.

"We have been in contact with Microsoft today. They claim that their intent is to only filter out the known bad hostnames in each seized [primary] domain, while continuing to allow the good hostnames to resolve. However, this is not happening. Apparently, the Microsoft infrastructure is not able to handle the billions of queries from our customers. Millions of innocent users are experiencing outages of their services because of Microsoft's attempt to remediate hostnames associated with a few bad actors. Had Microsoft contacted us, we could and would have taken immediate action. Microsoft now claims that it just wants to get us to clean up our act, but its draconian actions have affected millions of innocent Internet users.

"Vitalwerks and No?-IP have a very strict abuse policy. Our abuse team is constantly working to keep the No-?IP system domains free of spam and malicious activity. We use sophisticated filters, and we scan our network daily for signs of malicious activity. Even with such precautions, our free dynamic DNS service does occasionally fall prey to cyber scammers, spammers, and malware distributors. But this heavy-handed action by Microsoft benefits no one. We will do our best to resolve this problem quickly."

Well, this generated a huge amount of flak. And again, I mean, it's difficult to understand what happened. I read elsewhere that Microsoft was intending to somehow deploy their Azure service in, like, some dynamic domain filtering fashion. It sounds like it just completely collapsed; that it couldn't handle the demand; that they were, as a consequence, unable to filter. Instead, they just blocked them all.

And so by Thursday, the day before the Fourth of July, so four days later, the updated posting from No-IP says: "We would like to give you an update and announce that ALL of the 23 domains that were seized by Microsoft on June 30 are now back under our control. Please realize that it may take up to 24 hours for the DNS to fully propagate, but everything should be functioning within the next day. One of the domains, noip.me, took longer to get back online, but it should be fully restored within the next day. Is your service back up? Please send us a tweet and let us know." And then they sign off saying: "We're so sorry for the inconvenience that this takedown has caused our customers. Thank you so much for the support and for sticking with us," blah blah blah.

So this was just a - what Microsoft explained is that there were some botnets using this facility, and they were determined to take control of the DNS in order to get control of these botnets. And it just looks like, I mean, I guess we'll never really know behind the scenes what was going on, whether it was too much domain activity for Microsoft to dynamically filter, whether this was just deliberately done, we want all of these and we'll remediate them ourselves, who knows. But I hope that everybody learned a lesson from this because this was just - this was a real mistake for Microsoft to essentially kill a huge number of domains for a period of up to four days.

FR. ROBERT: Yesterday I had a little chitchat with one of my Microsoft friends. And he was basically saying that this whole thing started off with good intent. They intended this to be like the takedown of a botnet. But what was the problem is they had legitimate traffic mixed in with the malware traffic. And so they thought that they could drop this on Azure; they could run, as you said, the dynamic filtering; and they would allow through the legitimate traffic while blocking the malware traffic. And especially the command-and-control traffic for the botnets.

They ran into two problems. The first is just the sheer volume of traffic. They weren't expecting that much from a single ISP. The second thing is they started realizing that there are a lot of these domains that are going through No-IP that may be compromised and are therefore acting as either command-and-control or as attack vectors, but are still legitimate sites.

Steve: So they're both. They're both bad and good.

FR. ROBERT: They're both bad and good. And so the way they had set it up was, well, we're just going to find the bad domains, and we're going to kill them. And it was, well, that's got some bad traffic, but most of it's good traffic, so what do we do with this one?

Now, the issue I have with this is they were after two people. They were after Mohamed Benabdellah and Naser Al Mutairi, two black hat hackers who have kind of made a name for themselves for trading botnets. And they wanted the command-and-control for the

botnets that those two men had created. And as a result they took down an entire ISP.

Now, you can say that Microsoft had good intentions, and actually I'm willing to give them the benefit of the doubt and say they weren't just trying to be idiots. I think it was just really, really stupid. It was really ham-handed. They didn't do correct testing before they made the switchover. But what worries me is that they were able to use a security excuse to essentially seize the property of a competitor. They compete against No-IP. They run Azure. That, for me, that's a huge issue. I mean, it would be like Ford saying, you know, GM really screwed up with the ignition lock, so we're just going to take over their business for a while until we get them back on their feet. We don't do that.

Steve: Yeah. Well, and it's clearly a function of the court system. Some judge was confronted by a Microsoft attorney who I'm sure was very convincing and did a song and dance and got a court order. And you can do whatever you want to in this country with a court order.

FR. ROBERT: He said basically, oh, yeah, yeah, so here's the good thing. You're going to look awesome because you signed this order. And all of these customers are going to get all of the service that they expected, but we're going to kill a botnet. And the judge probably said, oh, yeah, that sounds good. Yeah, I'll sign off on that.

Steve: Yeah, you're right. And I'm sure they painted it to the court the way they expected it in the best case to work, which of course was not at all what happened. So it does sound like there were some lessons learned. And I hope observers learn so that they don't have to follow in Microsoft's footsteps. And I hope Microsoft keeps the people around who made the mistake rather than having them leave because now these people have learned a valuable lesson.

FR. ROBERT: Well, we hope. You can only hope.

Steve: So Oracle has made a sort of a strange announcement. And I don't know how to read this because they announced the end of their Java support for Windows XP, and it was picked up in the news and reported that way. ZDNet reported the regularly scheduled quarterly security updates for Java, the next one, which is set for July 15th, will not include updates for Windows XP, which is now formally unsupported by Oracle.

But then, when asked directly what Oracle plans to do, I mean, like, okay, well, what does that mean, ZDNet got this direct response from Henrik Stahl, who's the VP of product management for Java. And Henrik wrote: "As you know, Microsoft no longer supports Windows XP and recommend their users to upgrade to more recent versions in order to maintain a stable and secure environment. Oracle makes the same recommendation to our users running Java on Windows" - actually, one would argue, just don't use Java; but, okay, they're not going to recommend that - "and also has a standing recommendation that users stay current with the most recent Java security baseline" - okay, we're all for that - "currently available for the public for Java 7 and 8."

He writes: "There are a few compatibility issues with Java 8 on Windows XP, since it is not an officially supported configuration. We are looking at ways to resolve these." Okay, well, that's strange because they're saying they're not going to support it, but we're going to resolve them. And then he finishes, saying: "For now, we will keep Java users on Windows XP secure by updating them to the most recent Java 7 security update on an ongoing basis." Which seems to say they're going to continue updating Java 7 for XP.

"Java 7 users," he writes, "on more recent Windows versions can choose between Java 7

and 8 and, depending on their choice, will be kept up to date with the most recent Java 7 or 8 security update, respectively." So there he sounds like he's saying 8 won't be supported on XP, but Java 7 will on an ongoing basis. But then he says that they're looking at resolving the current compatibility problems with Java 8 on XP. So I don't think they know what they're going to do or what's going on.

And in looking, in digging deeper into this, I found, I mean, there's still a massive install base of XP. It's like a third of the Windows on the Internet. And so people are not leaving [XP], despite all the pressure on them to do so. So now we've got Java, which of course has been one of the largest security disasters in the history of the Internet, is Java. And Oracle's saying, well, we don't want to support it on XP anymore. The problem is many of these XP users are corporations which are unable to move because of compatibility issues, or are just unwilling to move, but who also use Java. And so here Oracle is saying, eh, we don't really want to support it anymore, but it looks like we don't have any choice.

So I really think there's a lesson that the industry as a whole is learning, which is you just can't force people to upgrade for your convenience. There is inertia, and there is a reluctance to leave something which is working. And it's funny because, against all of this, I was just reading, Mary Jo was tweeting the news about Windows 9 and that it looks like Windows 9 is essentially a complete capitulation to the understanding that 8 has been an unmitigated disaster and failure. They're going back to the traditional desktop. And so I'm delighted because I was thinking I'd be going from XP to 7. There's no way I was going to go to 8. Now I can just skip over both and probably go directly to 9.

FR. ROBERT: Which actually works because I've been telling people this. I run 8 on all my production machines. I don't use the Metro or the modern UI at all. But what I do like is what they did to the kernel. The kernel in Windows 8 is so much faster and so much more stable for some of the high resource usage software that I have, all my video editing software. So if they could drop the Windows 7 UI on top of Windows 8, I'm sold, absolutely. That's a no-brainer.

Steve: Nice. I think that's what we're going to get. Apparently 9 is a return to the look of 7. And they'll make it available. And as I understand it, what Mary Jo wrote was that they understood what they tried to force on the desktop was the disaster. Having it on the phone is fine. So the idea will be they recognize they're not going to succeed in pushing desktop users to the Metro interface. We want what we've got because it's effective. So 9 will be, exactly as you say, it'll be the 8 kernel with a traditional Windows desktop. And there will be Metro available. So it'll basically be both. And so on tablets and on the phone, where Metro makes sense, you'll have Windows 9 with that touch interface. Where it doesn't make sense, we'll have something that looks much more like we're used to.

FR. ROBERT: Now, Steve, as far as Windows XP and Oracle with Java is concerned, as you were saying, as you were explaining the story, it actually made sense to me because I've seen this from Oracle before. They're doing this with Solaris. So anything that they got from Sun, from that acquisition, that's free.

Steve: Ah.

FR. ROBERT: They don't like it because they can't monetize it really, really well. And Java is definitely one of those things where everybody expects it, but they make no real money off of it unless they're litigating. And they've got a history of doing this. This is a sort of a soft end-of-life. They don't want to tell people we're no longer supporting it

because they know there's going to be a backlash. But as they did with Solaris, and as they're now doing with Java for XP, they're saying you may run into problems. You probably just want to stop using it and go to something else. And I think that's just - that's their MO. That's what they do. They don't ever like being the bad guy, even though they look like it.

Steve: And doesn't that sound exactly like what the TrueCrypt guys did.

FR. ROBERT: Absolutely.

Steve: I mean, it's like, okay, don't use this anymore. We're not doing anything more with it, so you probably want to go to somewhere else. Which is not to say that there's anything wrong with it. But we're not going to fix things in the future.

FR. ROBERT: It's the same kind of language. There might be something wrong with it. We don't know. We're not going to invest any more energy into it.

Steve: Yeah.

FR. ROBERT: Maybe Oracle learned from TrueCrypt.

Steve: So, interesting piece of news that I got a lot of Twitter traffic about, which was a crypto weakness that was discovered in smart LED light bulbs. This was - it launched itself as a Kickstarter project, massively popular. I didn't write down what the numbers were, but I remember reading that they, like, raised way more than 10 times as much as they were looking for. It was like LIFX, I think is the - I can't remember even the name of the light bulb. And for me it really doesn't matter because I don't mean to pound on these guys. The nice thing is they responded very quickly to this. But this is, I think, a really useful cautionary tale about what's happening with the Internet of Things. And it connects back to why I was very glad in Apple's announcement recently that they're getting into the home automation market because I trust Apple to do this right. Apple will not make stupid mistakes. Even their first release will not have problems like this did. And I'm - there are a couple takeaways from this.

So here's the story. It's a very cool concept. The idea was that these would be LED screw-in light bulbs. And Leo's talked about them, where you can change the color temperature and so forth. And they would be in a mesh network. So the light bulbs, if you, like, have them many in the same room or scattered around your house, they would be talking to each other. So, for example, light bulbs far away from your router don't have to have a direct connection to your router. They can talk to the next nearest light bulb, that'll talk to the next nearest one, to the next nearest one, to the next, essentially forming a chain in order to all be connected together. And again, they talked about how easy this is to set up. Anytime something is really easy to set up, you have to ask yourself, okay, how is it working? How is it that I can screw in new light bulbs, and they're just on the WiFi network? How did that magic happen?

FR. ROBERT: I can tell you.

Steve: Yeah. Because if it's too easy, then you've got to wonder about the security. So what the engineers did was - or some hackers. Some hackers took some of these apart, and they found the standard debugging pin, the so-called JTAG pins, which allows access to the memory. They found the pins on the processors in the smart light bulbs, dumped out the memory, reverse-engineered the microcode, and found a static AES key. Now, I'm sure in the, oh, we've got military-grade security nonsense that was part of this, they

were saying, you know, AES 256-bit encryption, military grade, you know, we see this all the time in systems that are not secure. The problem was they burned the fixed static AES key into the firmware, the same one in every light bulb. So the instant the engineers, the hackers saw this, they were able to decrypt the traffic moving between the light bulbs, and that exposed the user's WiFi password. So the WiFi password was encrypted beautifully, but using a fixed known key.

Now, the danger is a bad guy knowing this could stand outside your house and easily participate in the mesh and with any, you know, the moment there's traffic between the light bulbs, decrypt it, and the user's WiFi password will be there. So this was the problem. Again, to this company's credit, they immediately strengthened their security so that it wasn't this bad. But this is the problem. As I was reading this I was thinking, one of the things here is that, as far as I know, there isn't, like, an RFC establishing a secure protocol for this kind of application. We've got all these secure protocols for doing all the kinds of common things we want. Well, this is still uncommon. And what we need is security people to establish a protocol for how to do this securely. And then other companies can simply adopt that protocol.

And as we do, as all the companies do who are on the Internet now, they're using well-established, very secure, pounded-on, bulletproof protocols. But we don't have anything like that for the Internet of Things. And so companies like this are just making stuff up. They're saying, well, you know, we're going to solve the problem because there is no RFC yet for it. Well, we need one.

FR. ROBERT: And this is just an example of security through obscurity. They figured, well, yeah, okay, we're using a static key, but we're going to bake it into a chip that no one will have access to. They won't be able to read it, and it'll be fine. And any security expert worth his salt would have sat next to them and said, "You know you can't ever assume that anything you bake into an IC is going to stay hidden. You know that; right?"

Steve: Right.

FR. ROBERT: And this is easier to hack than - remember WEP? Like WEP-64, WEP-128?

Steve: Oh, yeah.

FR. ROBERT: All you would have to do is send enough packets where you could decipher what the key was. This is even simpler. Once you decipher what any of the keys, once you figure out what that static is, any packet you could figure out what the key is because you've got the cipher key.

Steve: Yup.

FR. ROBERT: Whoa. I think we should design all our security like that.

Steve: Yeah, well, and notice that this was easy for them to do. I mean, AES, you can get off-the-shelf code for that cipher. So this was trivial for them to do. All they had to do, well, I mean, solving this securely would be tough when you have access to the microcode in the chip because I was going to say, if you did an ephemeral Diffie-Hellman exchange, so you're dynamically establishing a key, the problem is that, if you know, if you're able to spoof being a light bulb, and a new light bulb entering the mesh is going to receive the WiFi network's password, then I don't see how you convincingly protect this.

So I mentioned some takeaways. The takeaways are it is imperative, in my opinion, that

your Internet of Things devices be on their own network, their own WiFi network. We're now seeing routers that have a so-called "guest network feature." And if yours doesn't, get a second router and set it up with its own password that your hardware devices can talk to. My point is I just don't think it's safe at this point. We're like in the Wild West where you're going to get arrows in your back. You don't want hardware, I mean, like Nest and Insteon, and everyone wants to be on WiFi now. Give that stuff its own network. Routers are no longer expensive. As we mentioned, they're now commonly having multiple passwords. Keep that stuff off of your internal network. Unless people do that, I mean, the people hearing this podcast probably will. Most people won't. They'll have one network. They'll screw the light bulbs in. And there'll be trouble.

FR. ROBERT: And then the light bulbs will screw up their network.

Steve: I was just going to - I was tripping over that.

FR. ROBERT: You were going there, I saw that.

Steve: I was going there, and I said, uh, uh, do I want to say "And they'll be screwed?"

FR. ROBERT: I have no problems with saying things that will get me fired. It's cool. Don't worry about it. Steve, what's going on in India?

Steve: So this is sort of interesting because I knew something was up days ago. I guess it was on, was it the 2nd? Ever since the whole certificate revocation drama, I've had my eye on Chrome's CRLSets. The certificate that was being used for revoked.grc.com was manually added to the header of Chrome's CRLSet in order to block that certificate, which of course is cheating. So I changed the certificate. So it's no longer blocked and not revoked because, as we know, Chrome doesn't actually check for certificate revocation, even though it is a revoked certificate. I wrote some monitoring code, which continually polls in the same way that all of Google's Chrome browsers do all over the Internet, it polls the master server for CRLSets and notifies me the instant that changes.

And it was a few days ago I was actually - I was in the process of getting ready to run out for my semi-annual dental cleaning. And my monitor announced a change to the header, meaning that something had just been changed in these manually added certificates. And because I was running out the door, I didn't have time to check to see whether this was essentially the hash of my current certificate, meaning that they were now blocking the new one, which they haven't bothered to do for months. So I posted to the newsgroups at GRC. I said, hey, guys, Chrome's CRLSet header just changed. You may know before I get back whether they're now blocking my newly revoked certificate.

It turns out that's not what happened. What happened was Google found that a certificate authority in India was issuing certificates for some of their subdomains. And of course that's not kosher at all. Really the only reason you do that is you want to be able to undetectably intercept traffic to Google. So this was a certificate authority issuing against Google's policies, because they have no right to issue Google certificates, Google certificates. And they were found in use in the wild.

FR. ROBERT: Wow.

Steve: So this update immediately shut down Chrome's acceptance of the certificate on Windows because the other part of this is, only Microsoft was ever trusting this Indian certificate authority. It wasn't in the security suite which Firefox uses. It wasn't in the Mozilla security stack, nor in Android, nor in Apple. So it was only Microsoft that had this

Indian CA in their root store. And so it was only - so this CRLSet update would only block Chrome from accepting the certificate on Windows because Windows wouldn't know not to. And this gave Chrome the ability not to, independent of Windows, which is exactly the way it was working, for example, with my revoked.grc.com domain.

And of course Google notified Microsoft, notified the Indian CA that, I guess it was an intermediate that they had issued was being used to mint the certificates, rather than theirs directly. So they gave an intermediate certificate to somebody who was using it irresponsibly to mint Google subdomain certs. So a little bit of drama on...

FR. ROBERT: Wait. Is a certificate issuer actually allowed to do that? Can you grant that power to an intermediary?

Steve: Yeah.

FR. ROBERT: I thought that was against the terms.

Steve: In fact, well, you're able to specify how long the path is. And intermediate certificates are now being used more commonly. If you look at, like, all of the standard big CAs are not - are using a certificate signed by their root to issue. And some intermediates are not able to re-sign, and some are. So this one had that power.

FR. ROBERT: Wow. Okay. I'm still kind of scrolling a little bit here. The only reason why they would do that, I mean, they must have known that they were issuing certificates that were going to be used as vectors for man-in-the-middle attacks; right? I mean, that's the only reason why you would issue certificates that don't belong to you.

Steve: Well, yes. And the time that we have seen this before is when an intermediate certificate was in some appliance that was being used to mint certificates on the fly in order to do SSL/TLS interception. So we talked about this, oh, maybe about three or four months ago. There was another intermediate certificate that was found being used in an appliance that was able to synthesize certs on the fly in order to essentially sign any, make a certificate for any domain that you were going to that would be transparently accepted. As long as the signer of that intermediate was trusted by your root in your OS, no alarms, no dialogues, nothing would be brought up.

The normal way you do this is you're in a company that is doing this, and you have to add the filtering cert to your own root store on your laptop. So, for example, to get on the network you need to accept this certificate. And what that does is that allows the appliance that's between you, while you're in your corporate environment and the Internet, it allows you to trust the certificates that the appliance is synthesizing on the fly. That's the right way to do it. But it is not transparent because you first have to add that appliance's certificate to your local store. If, however, the appliance somehow is able to get an intermediate certificate which you trust, then it's completely transparent. Dangerous as all get-out because it can do anything. It's literally able to sign anything.

FR. ROBERT: So what's the takeaway here, Steve? I mean, Google has issued a cert revocation, which there's something ironic about Google issuing a cert revocation. We don't have to get into that. But people who know about the story, they'd be, yeah, it's a little strange. Now that they've issued that revocation, are we all okay? Or do we still have to worry about these certs?

Steve: Well, they've notified Microsoft, and they've only protected Chrome. And I know from my own experience that it takes days, actually, for the CRLSet to get updated. You

have to be using Chrome. And there's some length of time. When they issued the certificate for revoke.grc, it was strange. It took, like, in some cases two or three days before Chrome recognized that. So it's not like it's an instantaneous update by all means. However, it won't be until Microsoft updates their certificate store that IE, I guess it's only IE because I was going to say, I think Opera is no longer using IE. Chrome is, but they've got their CRLSet. And Mozilla with Firefox is bringing their own store along, and they never trusted that Indian CA.

So I think it's actually only IE on Windows is now vulnerable until Microsoft responds. And I imagine that this is probably too soon to have been part of this Patch Tuesday. In fact, I looked at all those patches. There was nothing about an update or blocking of this. So we may see something coming out in the next couple of days for Windows, an emergency patch to remove this from the trust store of Windows.

FR. ROBERT: And in the meantime, if you do use Google, just make sure you're not ever using IE. Although I think our listeners know that already.

Steve: Right. And I would say, again, just don't use IE at all.

FR. ROBERT: I was trying to be more diplomatic than that, Steve.

Steve: That ought to be your browser of last resort. You know, use it when you have to run Windows Update or something that demands - or Silverlight or something. But otherwise, no.

FR. ROBERT: I find myself using it a lot, whenever I get a new computer, to download Firefox and Chrome.

Steve: It's the bootstrap browser.

FR. ROBERT: Basically that's what it is.

Steve: Yes, to bootstrap yourself onto the Internet.

FR. ROBERT: It's what I use before I get the real Internet.

Steve: Right.

FR. ROBERT: Okay. So can you give us something to cleanse the palate? Get away from horribly acting CAs?

Steve: So of course with you last week, we talked about the project that I was going to launch on the podcast here to survey the current state of secure cloud storage solutions. At that time, I was aware that there was a Wikipedia page for hard drive encryption. And someone tweeted that to me, and it's like, yeah, yeah, yeah, I've seen that. But I did not know there was a Wikipedia page for comparison of online backup services. And it is stunning. I mean, it's amazing. So that's sort of thrown me off. I'm not sure what I'm going to do. It's still worth looking very closely, because that's the only way it's useful to look, at the best of these systems.

But for anyone who hasn't seen it yet, you can probably just google the phrase, because this is in the URL, "comparison of online backup services." I'm sure that'll take you to Wikipedia. And it's an amazing piece of work. I mean, there's stuff there no one has ever heard of. As it was, I was getting people tweeting, oh, what about this, and what about

my aunt's right shoe backup and, I mean, like, crazy things. It's like, okay, well, I don't know about that. So for next week, which will be the first of these, I'll take a look at that page. Basically I'll figure out what to do in light of that Wikipedia. There's no point in us recreating that. It already exists.

But what I want to do is what we really specialize in on the podcast, which is really, really drill down. And whereas, for example, on the Wikipedia page it'll say "secure key management," for example, I don't even know if it has a column for that, and it says yes or no. Well, we need to know more about it than that. I need to know more about it in order to trust it. So I'll be looking closely enough at these things to be able to explain exactly what this means. But there's just no point in recreating this beautiful piece of work that exists on Wikipedia and is being constantly maintained and updated. So for everyone who's interested, Comparison of Online Backup Services," an amazing page on Wikipedia.

FR. ROBERT: Yeah, it's actually, I mean, I'm looking at it right now. The amount of information they have and the breakouts that they did so that you can find all the minutiae of each service to compare them against each other, this was a lot of work. This was some people really put in the time to make sure that you had at least all the bold points. But as you said, bold points don't translate into actual service until you start looking at that minutiae.

Steve: Right. And this is also - that page, well, again, it's more work than I can possibly put in. And it's not the kind of thing that you just create overnight. I'm sure this has evolved over time with new lines being added of services and new columns being added of features. And, oh, my god, when a new feature column gets added, they've got to go back through every single one of those services and figure out whether to say yes, no, or we're not sure, or is it red or green or yellow and so forth. So a huge amount of work.

I made a comment last week that I wanted to correct. We were talking about Silent Circle, and I think in the context of the Blackphone, maybe.

FR. ROBERT: Oh, yes, yes.

Steve: And I thought that was Moxie Marlinspike. I misspoke. It's Phil Zimmermann who is the guy behind Silent Circle. So thanks for whoever tweeted the update. And I wanted to just mention a new maybe good sci-fi-ish series premiering tomorrow, on Wednesday, on CBS. Apparently Spielberg is involved somehow, probably as an executive producer, which doesn't mean that much. But this is the one that stars Halle Berry called "Extant," E-X-T-A-N-T. And there's sort of a creepy morphing of "extant" into "extinct" that happens. And all we know, and this is not a spoiler because everyone knows this who knows anything about the series, that she spends a year on a space station floating around in zero G and somehow comes back pregnant.

FR. ROBERT: As you do.

Steve: So we're not sure - as one might, you know, when one is, yeah, being visited by strange extraterrestrials.

FR. ROBERT: Steve, I like your entertainment choices, but I have been fixated on one thing. This actually is always on the big screen in my lab. I've got a 42" monitor, and this has been playing for the last week. JammerB, if you can go to this, this is a real - this is from IP Viking. It's a real-time map of attacks that they're detecting on their network. And it is - I could stare at this forever. It is just fascinating because it gives you

everything from where the attack's coming from...

Steve: It's got like little phaser beams moving.

FR. ROBERT: Yeah. So it's like an animated attack map, and it tells you at the bottom what are the IP addresses that are attacking, what are they attacking. They'll tell you the services, the ports that are being attacked. And my favorite, I've been using this as sort of like my World Cup of cyberattacks. It will tell you the ranking of attack origins and attack targets. And strangely enough, the United States is typically at the top of both charts at all times.

Steve: It is. Okay. How do our listeners find this?

FR. ROBERT: This is at IPViking.com. I'll make sure that the link ends up in the show notes.

Steve: IPViking.com.

FR. ROBERT: Seriously, I watch this. This is my entertainment. I just sit there with a bag of popcorn.

Steve: Well, I mean, it looks like a map of Global Thermonuclear War, like from...

FR. ROBERT: This is "War Games."

Steve: Like from "War Games." I mean, oh, my god, look at that right now. Some huge thing just - I think we just lost California.

FR. ROBERT: Do you want to play a game?

Steve: We're still here, though, so - wow.

FR. ROBERT: Every once in a while you'll see this massive attack going from Australia to Taiwan or vice versa. It's almost as if they're trading barbs. And it's typically separated by about eight hours. But when it does that, because the amount of traffic, I mean, I think it's like 80 to 100Gb. It just becomes this huge swarm and then just circles around the target.

Steve: It is utterly mesmerizing. It is.

FR. ROBERT: We could just do this for the next Security Now!, just put this up.

Steve: Yeah, I got to write it down. Tell me again, IP...

FR. ROBERT: IPViking.com.

Steve: V-I-K-I-N-G dotcom. Okay. Very cool. Wow. Okay. I wanted to let people know, just as a public service, there's an interesting-looking project closing on Kickstarter. I don't know what I was doing when I saw down the right-hand side "Other news items." And it said, "Super PC that fits in your pocket." And I thought, what? I mean, once upon a time there were like little tiny PCs, and of course I have one, like the OQO was a neat little pocket size with a screen and a keyboard. But of course tablets have taken that over. So it piqued my curiosity. How can you have a PC, like a powerful PC now, in your

pocket? What would that be?

It turns out what this is is an interesting Kickstarter project which is very close to making its, I think it's looking for $100,000 in order to kick it over. It was at 98-something when I looked an hour ago, 67 hours to go. So just a couple days, little more than two days, probably time for people to hear this. I just wanted to let people know in case it was the kind of thing that they would be interested in. The idea is it's a 100% solid-state module that you carry between docks. So the dock has the power supply, the fan, all your I/O interconnect, plugged into your keyboard and screen and so forth. And what you hold is sort of like a deck of cards-shaped thing. And so you plug it into your dock at home, use your computer, then pull it out, take it to work, dock it there, and use it there.

So again, I'm sure it's not for everyone. But I'm sure there are some use cases where this would just be like the answer for people. Oh, and in fact I'm looking at the page that you brought up, and they're now over $100,000.

FR. ROBERT: They just crossed the goal. So you're going to get them, folks.

Steve: So on Kickstarter it's Tango, it's called the Tango PC. Tango Super PC, a desktop Windows PC in a cell phone. So it's about the size of a cell phone. It's got a dual-core AMD, I think it was two gig, maybe. So a strong GPU. I mean, it's not going to be some kickass gaming machine. But for someone who wants, like, that kind of portability, where you unplug it from a dock and you take it to a different dock, it might just be the ticket. So I just wanted to let our listeners know.

FR. ROBERT: Is it just like a NUC? Is that what we're looking at? It's a different format? The New Unit of Computing?

Steve: Oh, yes, I think it is. It is a form factor that I've not seen before. So I thought that was sort of an interesting idea.

FR. ROBERT: And who doesn't like the idea of a docked something? We all love docks. In fact, Jeffrey's showing me his NUC right now. This is the one that he's been playing around with. And there's something that's just kind of cool about having a whole lot of power in something that's very quiet and very portable.

Steve: Yeah.

FR. ROBERT: That's not a laptop. Go figure.

Steve: Okay. So there was a tweet this morning that I got a kick out of because this is the beginning. Someone named "bothyhead," B-O-T-H-Y-H-E-A-D, and he's just @bothyhead, at 4:38 a.m. this morning via Plume for Android, tweeted: @SGgrc I've just been playing with Ralf's SQRL client and his test site. I so hope this takes off. It's amazing. The world owes you one." So what this says is, obviously, SQRL is running. And it is the case that there are going to be endless squirrel jokes, I'm sure.

FR. ROBERT: Either SQRL is running, or this individual needs a little bit of help, it's one or the other.

Steve: I mentioned Ralf a couple weeks ago when I was talking about the AES-GCM cipher protocol and how it was actually in my interactions with Ralf, who is a German student who is doing his master's thesis on SQRL, and also implementing an Android client and a test server. He was concerned about the intellectual property rights of OCB,

which is the cipher suite I was going to use, the authenticated encrypted cipher suite. And he raised some good points. I changed the spec and wrote, spent a week writing in Portable C an implementation of AES-GCM so that all SQRL implementations would be able to have one that was free, public domain, and completely unrestricted, since I wasn't able to find one otherwise on the Internet.

He's got his client up and running. A whole bunch of people over in the GRC newsgroup, the SQRL newsgroup, have it up and running and have been sending him feedback, like with what version of Android and what platform and what tablet and so forth. So it's beginning to happen. So it's all I've been working on. I'm working on the reference Windows client and working as hard as I can to get to the protocol portion because I just want to ratify the protocol, which is at this point still pro forma until I have a chance to nail it down. But it is the case that the SQRL system works, and it's working. So just a nice little bit of good news.

FR. ROBERT: So for all the people who have been hounding and hounding Steve Gibson on whether SQRL will be up and running, hey, go play.

Steve: Yeah. I'm still not done with mine, but I'm getting close. I'm working on - I've got all of the file system stuff, all the UI is in place, and I'm moving forward as quickly as I can.

FR. ROBERT: What's been the most difficult part about that entire project?

Steve: It's really the UI. I mean, for example, mine's Windows hosted. And I want to make it so simple to use that anyone can use it. So you run the client, and it says, "I didn't find any SQRL identities on this system." Well, that means that it has to know where to look. Well, I decided by default they would be under the My Documents folder in a SQRL, S-Q-R-L, subdirectory because that's a good per-user place for them to be. And in a corporate setting, where people are roaming, their folder tends to find them of their documents. So that means their identity would find them.

But then some people want to be able to use SQRL in a portable mode, where they would have it on a USB stick, which they would bring to someone's computer. That means their identity, their SQRL identity needs to be there with the client, with the SQRL client. And then some people said, yeah, but there ought to be some way to override that so that for specific corporate environments we could specify where the identity's going to be. So I said, okay, then we'll use the current working directory also.

So now there's three places my SQRL client needs to look for identities. It needs to look at a SQRL subdirectory in whatever is the current My Documents folder. It needs to look in the client's own execution directory and in the current working directory, which by the way you're able to use shortcuts in order to, like, easily, through a UI, set that to be anything you want. So it'll be in any of those three places. So that's done. Except what if there's a name collision? Because now you've got different directories, and you could have people who have named their identities the same. So I have to handle that.

It's like all of the plumbing of this just, I mean, the cryptography is relatively simple. It's all of the UI and dealing with users and, like, when someone says I want to create a new identity, well, okay, wait a minute. If you understand that the whole point of this is you don't need an identity per site, you just need one identity, and SQRL creates anonymous identities for you per website. So there's also this making sure someone who you don't, you know, no RTFM, somebody who isn't up to speed, who's just like the family guru said, oh, go use SQRL, I need to make sure they sort of get a tutorial automatically. So it's all of that. I mean, that's where the time has gone, or is going at the moment,

although I'm nearing the end of that because I've solved all these problems.

FR. ROBERT: Fantastic. Steve, I saw this in the notes, and I'm actually curious about this, too. You had a user who contacted you about running SpinRite on a BitLocked drive.

Steve: Yes.

FR. ROBERT: It shouldn't matter; right?

Steve: It absolutely does not matter. I just wanted to make sure people knew that. Kevin Marken, I ran across this in the mailbag this morning when I was pulling our Q&A questions together. He says: "Hi, Steve. I listen to Security Now! every week. I have looked a bit but have now - I guess he meant "not" - found anything about how SpinRite will handle a Microsoft BitLocked drive. Can SpinRite work the same on a BitLocked drive as a non-BitLocked drive? Can I run SpinRite against a BitLocked drive?" Signed, Kevin. And the answer is yes. Just as with TrueCrypt, SpinRite will see the partition table, and that will tell it where the partition is, and that's all it needs. SpinRite does not care what your data is at all. SpinRite 7 will be able to do that optionally.

Some of the next-generation things I'm going to bring into 7 is I'm going to fully tackle the full file system recovery and file level recovery. So SpinRite 7 will be file system aware when it has access to the file system. But we've talked many times about, for example, SpinRite recovering TiVo drives, where TiVo is Linux on a PowerPC where the byte order is swapped. And believe me, it has no idea what's on the drive. But it recovers it anyway. So, similarly, SpinRite can recover BitLocked and TrueCrypt or any other hard drive encrypted drive because it doesn't care what's there. It just makes the drive readable again.

FR. ROBERT: It's so low level that it's not looking for data, it's just looking for sectors. It's looking for what's actually on the disk.

Steve: Exactly. Exactly.

FR. ROBERT: Steve, I am afraid that I sidetracked you so much that this Q&A episode has very little Q, probably not a whole lot of A. Shall we jump at least a little bit into this?

Steve: Yeah. I would say, I'm looking at the clock, why don't we - we have, like, 12 minutes before 3:00. I bet that'll still give us a really great podcast, and any questions we don't get to we'll cover either next week or the week after.

FR. ROBERT: We'll get to it eventually.

Steve: Yeah.

FR. ROBERT: All right. Lead us off.

Steve: Okay. So Shane Elliott tweeted on Twitter, he said: @SGgrc Hi, Steve. Big fan of Security Now!. Was wondering if you know a reliable shared hosting provider good for developers. I'm using Media Temple now, but I like to shop the competition every few years in tech to find possible alternatives. Thanks.

And actually I was glad I had you here, Padre, because I thought, I'll bet that you know. My traditional go-to hosting provider has always been DreamHost, who has been, like, in the old days was really good. I thought something happened to them. Maybe it was they

got acquired by somebody else? I don't remember what happened. But who do you like for, like, tech-level, developer-friendly, website hosting providers?

FR. ROBERT: The one that I've been using over the last couple of years, and it gets a lot of bad press for some early problems that they had with developers, but it's 1&1. It's not the cheapest provider, and it's not the most full-featured provider. But they've always been able to give me anything I want. And they're actually really, really quick when you need to change something because you're testing a new feature or a new site. I've probably been working with them for at least, what, eight years, nine years now? So, yeah, that's definitely on the top for No. 1, as far as developing individually.

If I was going to go a step up from that, depending on how much you want to pay, Rackspace and CenturyLink both have really, really good plans for high-end development, but you are going to pay for them. So there's a bunch of providers in between the two. But those are the ones that I've used in the past with really good results.

Steve: Great, thank you. I think that's exactly what I was hoping to get from you. I got a note from Ken in Pennsylvania, who wonders about tricking XP into five more years.

FR. ROBERT: Awww.

Steve: Yeah. He said: Hello, Steve. I've seen on the Internet that there is a registry modification that will allow the consumer version of Windows XP - Home, Media Center, or Professional - to continue to receive updates from Microsoft until 2019 by making the Windows Update website think that your copy of XP is the POS version of XP, aka POSReady. The .reg file contents are as follows. And then he gives the link. He said: I've just made the change on my test machine and another 60 or so updates showed up on the Windows Update website. No reboot required after making the registry change. What are your thoughts about this? I haven't seen any negative comments from people who've made the change.

Okay. So here's what happened. That prompted me to do what I had planned to do, which was to make the change on an XP SP3 machine that I just sort of use as my mail station. I've got a little electronic scale and a dual - a postage and a label printer. And I just - I only use it, it's like a little turnkey just for weighing and preparing postage for things. So this morning I turned it on. I ran updates. And it is XP SP3, and it's been continually updated all the time. So I got an MRT update. And I'm blanking on it. Microsoft...

FR. ROBERT: Oh, is it - not the Malicious Software Removal Tool, it's the...

Steve: Yeah, yeah, yeah, yeah, yeah.

FR. ROBERT: MSRT, right.

Steve: Yeah, yeah, yeah. MSRT. I just wrote it down, that's why I didn't see it. Yes. The MSRT update happened. And it wanted to update Security Essentials, which it then shut down. Security Essentials was running just fine. It did a complete scan on the machine, everything was happy, and then it said, oh, we have an update to Security Essentials.

FR. ROBERT: It shut it down.

Steve: Which it then, it deliberately turned it red and said no more. You are no longer

supported.

FR. ROBERT: That's Microsoft saying, oh, yeah, we notice that our software is running. We should turn that off.

Steve: Yes. It worked perfectly. Then it said, okay, no, we're no longer going to provide this service on XP. So then I added the registry tweak to say that this is an embedded system. And I received seven updates from 2009. So old updates. And it called it WEPOS, which is Windows Embedded for Point of Service, and POSReady, and nothing else. I then rebooted, tried it again, nothing. So I put the link on GRC's server in case anyone else wants to see if they can reproduce this. It's GRC.com/fivemoreyears.reg. So it's just a little tiny registry file that adds this one link to the registry.

My guess is that Ken did it before this Patch Tuesday, and it was still working. And that, with this Patch Tuesday, they said, okay, we're closing this. We're not going to allow people to update their Windows XPs for - this was supposed to go to 2019, thus five more years. But I think the jig is up. And in fact, when we first announced this a few months ago, I said to Leo, I was like, you know, this won't be hard for Microsoft to turn off. And if it becomes popular, they'll probably - all they have to do, I mean, they could still honor POS systems by looking more closely to see if it actually is. They were doing a very lazy test by simply looking for this one key being set in the registry and using that as the sole determiner of whether this was truly a Windows Embedded system or not. And so it was a simple little spoof that was also not long-lasting.

FR. ROBERT: Yeah, if I remember correctly, the registry key that they were looking at just had the actual version name. And as long as you had the right version name, it would accept you as a Windows POS system and would update you. But they can do everything, like, look at licensing. They could actually request the licensing key, at which point you won't be able to reproduce that unless you actually have a POS license.

Steve: Yes, yes. It's trivial for them to look more carefully, and all they had to do was update Windows Update so that it actually would look more carefully at what was going on.

FR. ROBERT: Yeah. Awww.

Steve: So Opher Banarie in Chatsworth, California had a couple questions. In fact, I think he's got four questions, so this will be a perfect wrap for, yeah, I have four questions in one question. So this will be a perfect wrap for this week's Q&A. He wanted to talk about the EFF's Open WiFi initiative that we talked about a couple weeks ago, where the EFF is promoting OpenWiFi.org and firmware which they'll be offering when they announce it, I think it's next month or later this month because I think now we're in next month from when we talked about this first. They'll be announcing this at a conference later in the month for a router that's as yet still unspecified, which allows you to make Open WiFi available in a secure fashion by having an open and a closed WiFi network from a single router.

So Opher says: Hello, Steve. The discussion in this week's Security Now! about Open WiFi brought up some questions you didn't answer. One of the questions in the Q&A came close, but let's start at the beginning. First, the idea of violating ISP terms of service by knowingly allowing others to use the bandwidth is scary. Have you heard anything from EFF about what they are doing on this front? Okay, so, no. And I agree. That's a problem because, as we know, ISPs often have in their terms of service your implicit and explicit, when you say yes I accept these terms of service, statement that this bandwidth is for you and your household only, and that you agree not to be giving

the bandwidth that you're getting from them under these terms of service to anyone else, not making it available.

So all the EFF has said there is they are encouraging ISPs to change that, to remove that limitation from terms of service. And the EFF does provide a list of, oh, I think about 10 relatively small ISPs. When I scanned that list, no big ones like Cox or Comcast or AT&T leaped out at me. They were FrogFarm and things no one's ever heard of before. So it's like, well, okay. It's good that those guys are on the list, but we need the people, the ISPs that people are actually using to be on the list.

Then he asks: You've often covered how corporations use network hardware to spoof secure sites in order to read the traffic. Can or will Open WiFi enable private individuals to do the same? And I would have to say yes. Anyone who is using - and this is true generally. I mean, it's true at Starbucks when you're using unencrypted wireless. Even though this system uses encryption, it is only - it uses an encrypted link. It is only encrypted to the router. So any users of this Open WiFi need to treat it with the same level of caution they would any Open WiFi, which is, unless you have an SSL/TLS/HTTPS secure tunnel connection to servers, then you can't trust it.

I mean, it's still - you have to assume that this is going to be open and nonencrypted and subject to sniffing, perhaps by the people who are offering the service. I mean, that's really the danger is that whoever is offering this open wireless connection could be watching everyone who uses it. So you absolutely have to treat it under that assumption.

Third question: Will this guest network be forced to use the same DNS service configured on the private side? I could imagine someone not wanting their kids to figure out how to escape the neighborhood, so they block DNS to Google Maps. Will I not be able to connect to Google Maps using their guest access? I don't know for sure. We'll have to wait to see how the firmware works. It is certainly possible for firmware to deliberately block and/or redirect DNS. That's easily done. But normally routers will allow you to configure your own DNS locally and will honor DNS traffic going to servers other than the ones that it's offering through its DHCP service. So I would guess that you could manually override DNS and not be forced to use theirs. But it is possible that they could force otherwise. I would just be surprised if they did, especially coming from the EFF, that is, with firmware coming from the EFF.

And lastly, in a crowded environment such as a large apartment building, he asks, if there's only one Open WiFi Access Point that everyone uses, wouldn't that impact the speed available to everyone connected? I don't suppose hardware designed for home use has much in the way of muscle specification. And this is - that was my favorite question, only because I forgot to mention this when we discussed it.

One of the cool features for an individual who wants to offer this is that this firmware has a bandwidth limiter built into the Open WiFi side. So you are able to partition bandwidth so that the user is able to use - the visiting side is able to use bandwidth that you're not using, but you get priority access, and you're able to set a minimum allocation that they are always able to get. And they get any that you're not using, but you have first dibs on the bandwidth. So that's a really slick feature of this firmware that I failed to mention, which I think is clearly important. For example, if you were in an apartment building with this Open WiFi router, you wouldn't want the whole building using your bandwidth and starving you of ever being able to get any. So this thing works on a priority basis where the owner of the router gets first access to the bandwidth, and the freeloaders, for lack of a better term, get what's left over.

FR. ROBERT: That's decent QoS, but I will say that it's very easy for you to go ahead and

prioritize the bandwidth for the owner of the router or the access point. However, there is a physics limit here. And that is, if you have too much RF in the air at the same time, there is no prioritization over RF energy. So if there are so many clients or so many APs operating at the exact same frequency, on the same channel, you will see a degradation. But that's easily fixed as long as you know how to fix WiFi, how to properly configure it.

Steve: Right. Well, and it's interesting, too, because we sort of forget that this is also Ethernet. And Ethernet is about packet collision. And we're all using Ethernet where packets collide, and they back off and retransmit, and of course WiFi is using a shared medium, too. In this case the air is the shared medium. And as you say, if everybody's on Channel 11, then packets are going to collide. And we do know the one thing Ethernet does not do well is, when you get packet saturation, it fails rather badly. That is, you end up - if two packets collide, and then they both back off random amounts, but when they try to retransmit they collide with either themselves again or other packets, then they back off and try again. And you end up with, like, your utilization really drops at some point. That's the one failing of a shared medium like Ethernet where it just uses packet collision and random back-off in order to manage contention for a single resource.

FR. ROBERT: I think us old networking guys remember the names of repeaters and hubs, before switches became cheap. And you'd see the decrease in bandwidth. And it's not a smooth decrease, depending on the number of users. It decreases exponentially because you reach that point where most of the traffic is collisions. And the same thing happens in the air.

Steve: Right, exactly.

FR. ROBERT: Steve Gibson from GRC.com. He is our security guru. Steve, it is so - it's such an honor to be able to sit with you and chat. I've been watching you for so many years. To actually be able to have the last two weeks to just chew the tech fat with you has been a dream come true.

Steve: Hey, I did not want to forget to have you tell our listeners, this podcast's listeners, about your podcasts that you do on the TWiT network because I got so many really great tweets from people who said, wow, last week's podcast with the Padre was great. I thought, let's use this as an opportunity to make sure they know where you are.

FR. ROBERT: Awww, thanks, Steve.

Steve: The rest of the time.

FR. ROBERT: Well, you're going to find me here a lot, actually, on the TWiT.tv network. On Mondays you find me doing This Week in Enterprise Tech at 2:30 Pacific. I talk about networking. I talk about datacenters. I talk about how we're connected around the world. It's actually - it's close to Security Now!. I'd say it's a cousin of Security Now! because we don't go as in-depth on the security topics, but we do geek out over a lot of hardware and services.

On Thursday you're going to find me twice, at 11:00 o'clock for Know How with Bryan Burnett. It's a DIY Maker show. We do a few fun things. In fact, this week I believe we're talking a little bit more about our remote control project. I'm going to explain how you use ports on your home router. And then we're actually going to dunk a computer into liquid and make it continue to work. And then at 1:30 you find me for Coding 101 with Shannon Morse. It's the entre into the world of the Code Monkey. And finally, on Fridays, 7:00 o'clock, it's the late-night show, Padre's Corner. Join us here at TWiT.tv.

Steve: Okay, now, now, now, okay. If you're going to dunk the machine in water...

FR. ROBERT: I didn't say water. I said liquid.

Steve: The fans - okay, liquid, liquid, yes. Then the fans are not going to spin, yet the liquid will still take the heat off the heat sinks. So...

FR. ROBERT: Well, it's even better than that. The liquid that we're going to use, the fans will still spin. They'll just spin very slowly. It's going to be fun.

Steve: Very cool. Very cool.

FR. ROBERT: It's going to be - we geek out.

Steve: So they become water pumps instead of air pumps.

FR. ROBERT: Exactly. They still move fluid, it's just the fluid they're moving is not air, it's this other fluid.

Steve: Nice. Nice.

FR. ROBERT: But Steve, again, such a pleasure. Steve Gibson, you find him at GRC.com. That's the place where you'll find SpinRite, which we talked about, I think, at length. SpinRite has been - it's been something that's saved my butt more than a few times. You need to find out if it will save yours. It's the world's greatest maintenance and recovery tool. Also you'll find ShieldsUP!. That's another one of his tools which I think I use that on a daily basis, as well.

You'll also find 16Kb versions of this episode, transcripts, and of course some great information about security and SQRL, soon to be released, as well as an active forum community discussing everything under the secure sun. If you have a question you can submit it to GRC.com/feedback. And maybe your question will be picked up for one of Security Now!'s future Q&A episodes, hopefully a Q&A episode that doesn't have me, so we can actually do some Q&A.

Steve: Leo and I have often run over like this. And I think our goal here is to provide a good, meaty podcast. And we did that today, so I have no problem with the fact that we didn't get more questions in. We got two hours' worth of really good tech stuff. So I think everyone will be happy. Thanks so much, Padre. This was great.

FR. ROBERT: Thanks, Steve. Now, you can also find all of the versions of this podcast at our show page here at TWiT.tv/sn and wherever fine podcasts are aggregated. You can also use our apps, or watch us live at live.twit.tv. We gather here, normally with Leo Laporte and Steve, Tuesdays, 1:00 p.m. Pacific, 4:00 p.m. Eastern, 2000 UTC. Again, live.twit.tv. And as long as you're watching live, jump into our chatroom at irc.twit.tv, and you can, well, talk to Leo and Steve. I'm Father Robert Ballecer in for Leo Laporte. Thanks, Steve, again. We'll see you next week on Security Now!.

Steve: Thanks so much.