SECURITY NOW!

Transcript of Episode #570

Listener Feedback #238

Description: Leo and I first catch up with the past week's security happenings, including Apple getting Stagefright and speculation as to whether Russia is trying to influence the U.S. presidential election. Microsoft battles and wins against U.S. privacy overreach. Grace Hopper, who coined the term "software bug," brilliantly demonstrates a nanosecond. We've got a bug-fix update to pfSense, a "doing it weird" look at the CUJO security appliance, a bunch of errata, a bit of miscellany, and a dozen notes and questions from our terrific listeners.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-570.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-570-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here, and we're going to talk about the latest security news, some good court decisions, some bad court decisions, and of course it's question-and-answer time. And we've got not one, not two, not 10, but 12 question from you, our audience. Steve will answer them in an amazing display of intelligence, acumen, and perspicacity, next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 570, recorded July 26, 2016: Your questions, Steve's answers, #238.

It's time for Security Now!, the show where we cover the security of your stuff. And this is the guy who does it, Mr. Steve Gibson from GRC.com. Great to see you again. Happy Tuesday, Steve.

Steve Gibson: Great to be back, Leo. We've got a bunch of interesting stuff. Not a ton of security news, so I did a dozen Q&A questions from our listeners, rather than our usual 10, sort of judging how long this is going to go. But we have the news of Apple getting Stagefright, essentially. The question of whether Russia is trying to influence the U.S. presidential election. Microsoft's battles and wins against U.S. privacy overreach.

Something so cool, and I thank one of my Twitter followers, or one of our podcast followers sent me a link to Grace Hopper, who of course famously coined the term "bug," that's where "bug" came from, brilliantly demonstrates a nanosecond, shows us a nanosecond, and also relates it to a microsecond, I think it is. Maybe it's millisecond. I think it's microsecond.

There's a bug fix update to pfSense. I don't really have a "doing it wrong," I have a "doing it weird." Which is a bizarre look at what a consumer security appliance known as CUJO does, the way it connects to a network. Just kind of unnerving. We have a bunch of errata, a sort of an embarrassing quantity, frankly, of errata. A little bit of miscellany, and then a dozen notes from our listeners.

Leo: Don't you be embarrassed, Steve. Everybody makes mistakes. The only measure is...

Steve: Well, these are good ones, so....

Leo: ...how quickly you correct them.

Steve: These are good ones. I won't step on it by giving it away. Our Picture of the Week is a snapshot from "Star Trek Beyond."

Leo: Oh, did you see it?

Steve: Oh, yeah, of course.

Leo: Of course.

Steve: On opening day. That's really important.

Leo: You are serious.

Steve: So I called it "an engaging action film set in our J.J. Abrams rebooted Star Trek universe."

Leo: Right.

Steve: And I think that's the way to frame it. I mean, it's an action film. I know that there are diehards that want more Roddenberry-esque meaning of the Federation, I mean, and there was homage to that. But mostly it was a lot of fighting, more than I needed. But I liked the movie, I mean, it was Star Trek. And lots of beautiful...

Leo: Aliens.

Steve: ...consoles, and I like the new crew.

Leo: You were going to say consoles. I said aliens. But they were both beautiful, yeah.

Steve: Yeah.

Leo: Good blinkin' lights, in other words.

Steve: And she really can't act very well, that blonde. It's like, oh, goodness. But, you know.

Leo: It's hard to act when you're wearing whatever that is on her head.

Steve: And actually, some of the lines they have to say, like, "Where is your home planet?" There's no way you can say "Where is your home planet" with, it's like...

Leo: Yeah, where are you looking in your inner life, your prior experience, to give that some depth? And who knows, I mean, yeah. Let's move on.

Steve: So Cisco's Talos security group we've talked about actually more and more recently. These guys are doing a great job, sort of in the same way that Google's team is, looking at, just in general, at things in the industry and giving them a onceover. And in this case, Tyler Bohan of Cisco Talos discovered five remote execution code vulnerabilities in various pieces of image-rendering code in OS X and also iOS, that is, in the latest versions before this very recent update. So he of course did responsible disclosure, informed Apple. They fixed it. And so once the patches were out and available, we got the news of that.

And so I called this a Stagefright bug because it's very reminiscent of Stagefright. As we know, of course, the Stagefright was dogging, or actually still, really, is a module dogging Android because it handles a lot of the media processing. And in the case of Android, when we first saw Stagefright happen, it was just receiving a multimedia SMS, an MMS message, that the image in the multimedia event would automatically be processed by the Stagefright module and, because there was a mistake there in the parsing of the image, it would allow a bad guy to essentially put their own code in with the image and get it to execute.

Well, Apple got hit by the same thing, and in a very similar way. One of the things we've been talking about the last few months is the difficulty of not making a mistake when you're coding an interpreter because an interpreter, the person coding an interpreter sort of assumes that what it's interpreting will be sane, that it's going to interpret valid input because why wouldn't it? Well, it turns out bad guys have exactly the opposite approach. They look for subtle mistakes in the interpretation path that they can take advantage of.

So in this case, and there was an interesting lesson here we'll get to at the end, the big mistake was in Apple's handling of, believe it or not, TIFF files, the Tagged Image File Format, which I'm having to tell people what TIFF stands for because unlike PNG, JPEG, and GIF or GIF, however you pronounce it, few people these days even see a TIFF file.

But in case one comes along, the code is still there for handling it.

So what Talos wrote I thought was just so good, I couldn't even paraphrase it without changing anything, so they said: "The Tagged Image File Format is a file format that's popular with graphic artists, photographers, and the publishing industry because of its ability to store images in a lossless format. TIFF was created to try to establish a common scanned image file format in the mid '80s. Cisco Talos has discovered a vulnerability in the way in which the Image I/O API parses and handles tiled TIFF image files. When rendered by applications that use the Image I/O API, a specially crafted TIFF image file can be used to create a heap-based buffer overflow and ultimately achieve remote code execution on vulnerable systems and devices.

"This vulnerability is especially concerning as it can be triggered in any application that makes use of the Apple Image I/O API when rendering tiled TIFF images. This means that an attacker could deliver a payload that successfully exploits this vulnerability using a wide range of potential attack vectors including iMessages, malicious web pages, MMS messages, or other malicious file attachments opened by any application that makes use of the Apple Image I/O API for rendering these types of files.

"Further, depending on the delivery method chosen by an attacker, this vulnerability is potentially exploitable through methods that do not require explicit user interaction since many applications, i.e., iMessage, automatically attempt to render images when they are received in their default configurations. As this vulnerability affects both OS X 10.11.5 and iOS 9.3.2 and is believed to be present in all previous versions" - again, this is old code - "the number of affected devices is significant."

So I sent out a note, I think it was just yesterday, as I was digging into this more because I was aware of 9.3.3, that is, iOS 9.3.3. But I think maybe only one of about, I don't know, I have, like, 12 iOS devices I manually updated, not a single one of them did that by itself. And so I wanted to alert everyone, you know, this would be a good time to go just check to see if there's an update available, and probably it'll say yes, we've got one. And it's like 50MB and requires the regular reboot and restart and so forth.

But for what it's worth, and for whatever reason, not one of my devices did this on its own. And for something like this, this doesn't want to hang out there for too long because, as we know, once bad guys realize there is an exploitable flaw, especially in iOS, where these generally are rare, and these guys didn't talk about what privilege the execution code would have, but it may very well be highly privileged execution, this could be bad.

So that's one of four. I won't go into the same detail with the other three, or at least the middle two. One was a similar problem, and again, a remote code execution in the OpenEXR file format. And actually there were two vulnerabilities there. And the TIFF only had one vulnerability. They also found a problem with the Digital Asset Exchange File Format, one vulnerability. And then also in the BMP, old standard Bitmap File Format.

And they wrote, they had a little short write-up for that. They said: "The BMP file format is both longstanding and has a fairly straightforward structure. The BMP file header contains information about the size, layout, and type of image. A vulnerability exists within the way that the height property of an image is handled. This can be exploited when a specially crafted BMP image file is saved, then opened, and part of the size information is manipulated."

Leo: Yeah, let me guess, it's more than 65,635 pixels tall or something; right?

Steve: Yeah. "The exploit leads to an out-of-bounds write, resulting in remote code execution when opened in any application using Apple Core Graphics API." So then, to paraphrase a little bit, they said: "Image files are an excellent vector for attacks since they can be easily distributed over web or email traffic without raising the suspicion of the recipient. These vulnerabilities are all the more dangerous because Apple Core Graphics API, Scene Kit, and Image I/O are used widely by software on the Apple OS X platform."

Then they said: "Organizations should patch software to the latest release in order to resolve these vulnerabilities. Additionally, organizations may wish to consider blocking" - and this is what I wanted to get to - "should consider blocking files at network gateways if the file is of a type that is never, or very rarely, going to be encountered within the legitimate business of the organization." For example, TIFF files. And that's significant because most companies could completely sail along with no TIFF files crossing their Internet to Intranet boundary. It just, you know, you just don't run across the file format any longer. Yet it's still supported for legacy reasons.

So sort of in the same way that we've switched our firewall concept from block what's bad and allow everything else, we've switched it around to drop everything, that is, block everything, and then selectively open the traffic that we know we want. It really does make some sense to consider where you have the ability to do content filtering to look at all the file formats that are around. And if you don't recognize them, you probably don't need that. And just say, eh, no.

I mean, the worst that could happen is an exception would have to be made in a specific instance. But, for example, anyone who had that kind of firewall up, who was blocking TIFF files because when's the last time you saw one go by, even if this were exploited in a zero-day fashion - and as far as anyone knows it has never been exploited, that is, Apple fixed this before this got out because Cisco Talos reported it responsibly. But the point would be that this is the kind of thing that, if you were preemptive, your corporation would be protected, even if it were found. And it generally is this legacy code that tends to bite people. And happily, things like Flash are becoming legacy as we move to HTML5.

So you want iOS 9.3.3 or OS X 10.11.6 or later because those releases have these things fixed. And again, I had to do it manually. So I would suggest iOS users who are concerned about this, just make sure that your devices are running the latest. Do you know, Leo, what the schedule is for this, like the updates?

Leo: So because Talos told Apple before they revealed it, Apple was able to update last week iOS, El Capitan, tvOS because it's in the Apple TV, and watchOS because it's in the Apple Watch. So those are all fixed. One of the four vulnerabilities is not patched on Mavericks and Yosemite, older versions of OS X. So that would be the only place you'd have to worry. But of course you want to make sure everybody updates. And not everybody does updates.

Steve: So is it an iOS user's responsibility...

Leo: Yes.

Steve: ...to go get the update? Apple doesn't notify you?

Leo: Yeah, you'll get a notification. But you may not get it right away. You may get - or you may just see the icon, the settings icon badged, that kind of thing. So everybody should do it.

Steve: I sometimes do run across my devices that have a fingerprint scanner, it'll say, oh, you have to enter your passcode after a restart. And I think, well, it worked yesterday. So apparently the device restarted itself at night for some reason, maybe to perform a silent upgrade like that.

Leo: I'm sure Apple has the capability. I don't think it does that. I don't think it did it with 9.3.3 on iOS, but maybe.

Steve: Yeah, it didn't on any of mine.

Leo: No, yeah.

Steve: So, okay, this is weird. And I have to cover it because we're Security Now!, and we're now in the second week of the two-week election conventions. And there's all this in the news about the claims that Russia may have been involved in hacking a Democratic National Convention staffer's email account and actively involving itself by selectively releasing some emails in the U.S. election outcome. Now, of course, I don't know whether that's true or not, so I can't comment on that. That's what's in the news.

Multiple analysts have confirmed that Russian state actors did penetrate the DNC email system, and also apparently some personal email accounts of DNC staffers, which, I guess because they were consultants, they were also using to conduct DNC business. The leaks came through WikiLeaks, yet Julian Assange is refusing to provide any attribution of the source either way. So he's saying, you know, the people who give us tips require and ask for anonymity. That's what we're providing. So there's no confirmation there. But Michael Isikof, who's a respected reporter, did report on this. And I have a picture in the show notes of the pop-up that this DNC staffer was apparently receiving for days.

Leo: And ignoring.

Steve: And ignoring, ignoring. So Michael writes in his coverage: "Just weeks after she started preparing opposition research files on Donald Trump's campaign chairman Paul Manafort last spring..."

Leo: By the way, I'm glad to see she's using Yahoo! as her home page.

Steve: I was going to say, you know. I have a hard time taking Yahoo! seriously. I don't know what it is. I just - I never have. It just always seemed like maybe a step above AOL, but as Mom calls it, AWOL.

Leo: But she obviously took a picture of this, so she must have wanted - she must have thought about it.

Steve: Well, okay. So Michael reports: "Democratic National Committee consultant Alexandra" - I guess this is Chulapa.

Leo: Chalupa, just like at Taco Bell, yeah.

Steve: And I was thinking, isn't that a hot sauce? But I think that's Cholula.

Leo: Yeah. But there is a Chalupa which is a Taco Bell treat, yeah.

Steve: Okay. So Alexandra Chalupa got an alarming message when she logged into her personal Yahoo! email account, and it reads "Important action required" as the headline in bold. Michael says: "...read a popup box from a Yahoo! security team that is informally known as 'the Paranoids.'" In this case, not so much. Or as they say, even if you're paranoid, it doesn't mean that they're not trying to get you. Then it continues in this dialogue: "We strongly suspect that your account has been the target of state-sponsored actors." Now, maybe she thought that meant...

Leo: A mime troupe.

Steve: ...Alan Alda, you know, or Clint Eastwood. I mean, I don't know, like she didn't understand...

Leo: We've got a Chinese mime troupe in here. We've got...

Steve: Yeah. So then...

Leo: But she did take - somebody took a picture of it. I mean, she must have taken a picture of it; right?

Steve: Well, she reported it. So get this, though: "Chalupa, who had been drafting memos and writing emails about Manafort's connection to pro-Russian political leaders in Ukraine, quickly" - and I put in my notes, quickly? - "quickly alerted top DNC officials, saying, and I guess this was her interview by Michaels: 'Since I started digging into Manafort' - get this - 'these messages have been a daily occurrence on my...'"

Leo: Oh.

Steve: How pesky. These pesky pop-ups warning me of state actors "have been a daily occurrence on my Yahoo! account despite changing my password often." And, I mean, I'm just gobsmacked. It's like, okay. Didn't go to Gmail. Didn't decide to get a DNC account. Just thought, well, I'll change my password. Oh, look, another one of these pesky "important action required" messages. Wow. And then "A Yahoo! spokesman said the pop-up warning to Chalupa 'appears to be one of our notifications' and said it was consistent with a policy announced by Yahoo! on its Tumblr page last December to notify customers when it has strong evidence of 'state-sponsored' cyberattacks. Bob Lord, the company's Chief Information Security Officer, wrote in that Tumblr post: 'Rest assured, we only send these notifications of suspected attacks by state-sponsored actors when we have a high degree of confidence.'"

Leo: I think Google does this, too. These are not - this is not just Yahoo! doing this.

Steve: Yeah. I mean, what you'd like is your account is locked forever. Go find a real email system. Anyway, who knows how they got in or how they managed not to get shaken off by her frequent password changes. But the idea that she was getting these daily and continued writing emails and memos in the face of notifications indicating that state-sponsored actors were hacking her account, again, we...

Leo: What should she have done? I mean, you know, what should she have done? Called the FBI?

Steve: I would say, I mean, okay. So she should have immediately reported it to the DNC, and they should have said, okay, stop using Yahoo! Mail. We'll set up an account for you at the DNC to use for DNC-related business.

Leo: But the DNC was hacked, too.

Steve: Yeah, that's not good.

Leo: So it doesn't - I wouldn't blame her too much. I mean, normal people, when faced with something like this, I mean, I don't think...

Steve: And it did, Leo, it had a very clear little X in the upper right-hand corner. Just like that...

Leo: Close this because it could be wrong.

Steve: ...that pesky notice goes away.

Leo: Right, could be wrong. I mean, it doesn't - it's not definitive.

Steve: So we didn't mention here, although you did cover it over the weekend, Microsoft's successful outcome with the New York-based or located Second Circuit Court of Appeals in this issue of our U.S. domestic law enforcement back in 2014 issuing a subpoena or a warrant for them to provide information that was stored out of the U.S. in Ireland. So this is the Second Circuit Court of Appeals sided with Microsoft in this case over whether the U.S. government could force the tech giant, Microsoft, and other companies to hand over customer emails stored overseas. So this appellate decision "reverses the original 2014 court order requiring Microsoft to turn over email which was stored in a server in Ireland."

So that decision happened two years ago. Microsoft said, "We're going to appeal it," and they did, and they won on appeal. And this was a narcotics case that these emails were believed to be connected to. A judge, Susan Carney, with the Second Court of Appeals, "found that the federal Stored Communications Act only applies to data stored in the United States, and thus cannot be used to force a company to produce information from servers outside the country."

So now, with their original warrant invalidated, the government must proceed through a much lengthier process to set up something called a "mutual legal assistance treaty" with the Irish government to obtain the data. However, Ireland filed a brief supporting Microsoft in this case, and they were joined by a bunch of other tech companies including Apple and Cisco. So it's not clear how cooperative Ireland is going to be in this. I mean, again, we're in this new place where we have strong crypto. We have a well-connected global Internet and pesky things like national boundaries, that never used to be a problem, we now have to deal with as this information flows freely across borders.

Leo: You don't, I don't think, in your show notes mention another court decision using the Computer Fraud and Abuse Act. It was against a company that was scraping Facebook information at the request of Facebook users. Oh, my gosh, this is a good one. Basically it says that, if you use somebody's website against their permission, you're violating the CFAA, and it's a felony.

Steve: Oh.

Leo: So if I put a big sign on here that says, "Hey, Steve Gibson, you may not use this website," and you do, you could have some serious consequences.

Steve: Wow.

Leo: Yeah. This is the California Court of Appeals.

Steve: Ninth Circuit, probably.

Leo: Yeah. I don't - yeah. So the issue was Power Ventures invited Facebook users to sign up. And then, as usual, give us access to your contacts, whatever. And they did that. And then Facebook sent them a cease-and-desist letter saying stop. And they continued doing it, and then of course the court battle ensued, and the judges say, no, no, once you got that cease-and-desist, you're violating the CFAA if you log in again.

Steve: Ooh.

Leo: Even if the user asks you to.

Steve: And so that got overturned?

Leo: No, no. That's the decision.

Steve: Oh, no.

Leo: It's a felony.

Steve: So it hasn't gone to appeal yet.

Leo: Well, it would have to go to the Supreme Court. The Ninth Circuit...

Steve: Oh, the Ninth Circuit upheld the CFAA?

Leo: Upheld.

Steve: Oh.

Leo: They sided with Facebook.

Steve: Wow. The courts...

Leo: They taketh, and then they giveth away.

Steve: That's the end of life as we know it.

Leo: Well, it could be. I mean, it just depends. I still think a judge can decide not to do it. But...

Steve: Yeah, this is bad precedent, though.

Leo: Oh, yeah. If Twitter says, hey, Nero, if you log into our site, you are now violating a federal law, and you could go to jail for a long, long time. The CFAA really is a blunt weapon and is often misused, as we know.

Steve: And in fact we were talking about it last week, that one of the problems - and I kind of got a little carried away talking about it, that some of these things start off being sharp; but in order to pass, they're deliberately blunted.

Leo: Right.

Steve: In order to not hurt various factions' feelings and in order to essentially buy their votes by weakening the law. And, wow. Ooh, boy.

Leo: You know, that's one of the things, unfortunately, I mean, we're going to have to operate on the, come November, on the state level and the local level because it's something neither presidential candidate knows anything, has any information at all about. But the President's not the person who's going to decide, it's Congress. So just keep these things in mind when it comes around November time, and vote for someone who does know.

Steve: Okay, now, the audio for this video is low, so you'll need to turn the audio up a little bit, Leo. And it's only two minutes. I think we should just play it into the podcast.

Leo: Absolutely.

Steve: And then I'll give everybody an easy way to find the link. This is Grace Hopper.

Leo: Admiral Hopper.

Steve: Who was - she was Navy; right? I think she was Navy.

Leo: Yeah, she was an admiral.

Steve: Yes, and an early programmer of the early mainframes, who coined the term "bug," who is explaining, brilliantly, the concept of a nanosecond, and then relating it to a microsecond.

[CLIP]

GRACE HOPPER: They started talking about circuits that acted in nanoseconds, billionths of a second. I didn't know what a billion was. I don't think most of those men downtown know what a billion is, either. And if you don't know what a billion is, how on earth do you know what a billionth is? I fussed and fumed. Finally, one morning, in total desperation, I called over to the engineering building, and I said, "Please cut off a nanosecond and send it over to me." And I've brought you some today.

Now, what I wanted, when I asked for a nanosecond, was I wanted a piece of wire which would represent the maximum distance that electricity could travel in a billionth of a second. And of course it wouldn't really be through wire. Be out in space, velocity of light. So if you start with the velocity of light and use your friendly computer, you'll discover that a nanosecond is 11.8 inches long, the maximum limiting distance that electricity can travel in a billionth of a second. Finally, at the end of about a week, I called back and said, "I need something to compare this to. Could I please have a microsecond? I've only got one microsecond, so I can't give you each one. Here's a microsecond.

Leo: She's pulling it out of a paper bag.

GRACE HOPPER: Nine hundred and eighty-four feet.

Leo: Wow.

GRACE HOPPER: I sometimes think we ought to hang one over every programmer's desk, or around their neck, so they know what they're throwing away when they throw away microseconds. Now, I hope you'll all get your nanoseconds. They're absolutely marvelous for explaining to wives and husbands and children and admirals and generals and people like that. An admiral wanted to know why it took so damn long to send a message via satellite. And I had to point out that between here and the satellite there were a very large number of nanoseconds. You see, you can explain these things. It's really very helpful. So be sure to get your nanoseconds.

[END CLIP]

Leo: I love it. She was handing them out.

Steve: Wasn't that great?

Leo: Oh. I wonder if that was a class or a - that's awesome. It looked like a bunch of...

Steve: So, and for those who couldn't see it, when she was talking about the distance to the satellite, she was taking this 11.8-inch wire and sort of moving it in steps, like putting it end to end to end to end, all the way up to where the satellite would be. And I also loved her explanation, I mean, that of course speed of light, as we know, nothing can move faster than, that it's the maximum limiting factor. That is, so here's this beautiful, straight, 11.8-inch piece of wire which she calls a nanosecond because it's a physical representation of a nanosecond of propagation distance.

Anyway, I didn't know I was going to have time to do this, so there's actually a Q&A note later where someone tweeted to me and said: "Steve, people are now messing with your bit.ly link formula. How hard is it to create a web page with links on it?" And it's like, oh, yeah. I have a web server. I could do that.

Leo: You can even do your own bit.ly.

Steve: Well, actually I have GRC.sc for shortcuts.

Leo: Yeah, there you go.

Steve: And it has always been my plan to do that. But we all know I'm a rather busy boy. So, and I don't want to do like a cheesy one. I want to do a, like, I want to do it once and forever solve the problem. So it would have a database on the backend and a nice UI that allows me to set things up and so forth.

Leo: Bit.ly has a white label version that you can do for free, by the way, if you want [crosstalk].

Steve: Well, I never got around to it.

Leo: Yeah.

Steve: So what I did was - and there are other advantages to have something browsable because, for example, people have asked, what about that, you know, where is your sci-fi reader guide? What was that site that you talked about for privacy information and so forth? And so I thought, okay. Years ago, in the newsgroups, people were posting links in, like, sort of just everywhere. And it was annoying people. So I created a group just for that purpose, which I called Link Farm. I don't know why. I just - I thought it was kind of funny. And so we now have a Link Farm page at GRC, GRC.com/linkfarm. And that will be the place from now on where I will just post links. I mean, the show notes always have them, but not everybody gets the show notes.

Leo: I like the little barn, the little red barn and silo you have there.

Steve: It's the Link Farm.

Leo: You should get an animated GIF with a little cow coming out of the door.

Steve: And there will be - links will be harvested, grown and harvested here. And again, I did have time this morning, I didn't expect I was going to, but I didn't have time to do anything other than that one link. But I will, as I can, populate this with some other often-requested things, and our listeners will know that at any time they can just go to GRC.com/linkfarm, and it'll be in most-recent-at-the-top easy format. And the other thing that allows, of course, is browsability. You can browse backwards in time. Now, I'm not going to go and repost all of the previous podcast links. But moving forward, we now have a place for those to go.

So whoever that was who suggested, "Uh, Steve" - and actually there was a little back and forth in Twitter, and he was saying, you know, just set up a 301 redirect or a bit of JavaScript or something. It's like, no, no, no. If I'm going to do it, I want to do it right. But as a consequence, nothing's happening. So now we have a simple page where I can put links. So no more bit.lys.

I did want to quickly note that pfSense has been updated to - now, I wrote 2.3.2. But I thought it was 2.3.3. Now I've confused myself. But it did just get an update. So anyone using pfSense, you can actually, if you just go to your main admin page, in the upper left-hand corner you'll see it checking for any updates to itself. And it will now say, yes, got one. And they fixed about 60 bugs, added eight features, and two to-do list items completed. I have a link in the show notes to everything that they changed. And this is what you want in a border router. This is what you want in a router whose security you're depending upon, people who actively care, who are fielding reports of any odd behavior and fixing it and then making it available. This is not what we have in our existing turnkey consumer blue box routers. That's why pfSense and I are getting along so well. So I just wanted to let people know there's an update.

And I mentioned this CUJO. This is a consumer appliance thing. And I'm sure you've seen it, Leo, or a picture of it. It's kind of an inverted bowl shape with a flat head, but then it's got two LED things that kind of look like eyes. And so it's supposed to be sort of a little friendly consumer appliance thing of some sort. And I'm not a fan of technologies that sort of try to - that claim to be able to do more than they probably can.

But someone sent me a note some time ago saying how can this intercept my network traffic if you just plug it into your router? And that's what you do. This little thing is an appliance. And the other thing I'm not a fan of is $9 a month for the privilege of this probably not being able to do that much for you. So there's that. But I was curious. How can it do anything?

Leo: You're not in between the person and the Internet.

Steve: Correct.

Leo: You're just sitting there as a peer.

Steve: Correct. Correct. And believe it or not...

Leo: Maybe the lights light up or something, eyes light up or something.

Steve: This is a consumer ARP attack.

Leo: Oh, nice. It's a man in the middle. Cujo in the middle. Oh, lord.

Steve: Completely breaks all the rules of how the network should work. So I found - I dug around, and I found their explanation. And our listeners will be able to read between the lines. They said - so the question. A knowledgeable person writes: "I have a highly customized router based on WRT, so I'm curious how it is going to get access to all the packets on my switched network without me having to make changes to my router settings."

Answer, they write: "The CUJO" - and that's, by the way, C-U-J-O - "appliance works in one of two modes." So they do have what they call their gateway mode. "Our Gateway mode, where you plug it into your router with a" - oh, I'm sorry, yeah, the gateways. "Our Gateway mode, where you plug it into your router with a single Ethernet cable; or our Bridge mode, where it sits between your modem/router and switch." For those who have those functions separate. Many people don't, especially if your router, for example, has WiFi. Then it can't get in between.

Get this. "Our Gateway mode works by intercepting packets via an ARP mechanism. This is how we achieve our 'simple plug and play' goal for the average Joe. Our Bridge mode works as you would expect. Because it sits in the middle of a modem/router and a switch, it's physically in the middle." Then they said: "Once the CUJO is logically or physically in the middle, we sample metadata from your network's connections," they said, "using NetFlow. The metadata is strictly src/dest IPs and ports, bandwidth, packet count and connection states."

They said: "We do NOT perform deep packet inspection as it is too intrusive and has a pretty big performance penalty for us." Actually, we know they don't perform deep packet inspection because they can't, because everything is encrypted these days, and there's no visibility into the packets going by. "These samples," they say, "are hashed and sent to the CUJO cloud over an encrypted channel. In the cloud is where we do the heavy lifting."

So essentially we have a deliberate ARP spoofing attack, meaning that when any device on your network sends a query out to get the MAC address of the router, that is, of the gateway, what's supposed to happen is that's a broadcast. Since the device on your network, a light bulb or an IOT device or whatever, has no idea where you are, it broadcasts it. Anybody can reply. This CUJO replies first, before your router is able to, claiming to be...

Leo: We have a video dramatization from the Stephen King movie of the same name.

Steve: Oh, "Cujo," right.

Leo: How you describe it, I think you'll see it's quite apt.

Steve: Oh, boy. So anyway, I'm not putting something on my network which is going to commandeer, by breaking the fundamental architecture of - oh, and, I mean, well, I've just stepped on myself - by breaking the fundamental architecture of the way networking works and the way, you know, ARP stands for Address Resolution Protocol. And imagine if you then add something else to your network that wants to do the same thing. Now they're, like, now you have three devices fighting for supremacy, and one of them is going to win, and not necessarily the same one every time. I mean, it's just an incredible kludge. So as I said, maybe, I mean...

Leo: [Crosstalk] going to win. Cujo. He's here. Scary St. Bernard. All right.

Steve: Was that an actress that we recognize?

Leo: Yes, from "E.T."

Steve: I thought so.

Leo: Yeah. By the way, that was an Indiegogo project originally, CUJO was. It was one of those crowd-sourced...

Steve: That's where I saw it. That's how it came on my radar was, yeah. Well, and again, maybe for a nave user - first of all, I'm not sure what monitoring hashed aggregate packet stream is doing.

Leo: They might see malware floating by or something like that.

Steve: Yeah. I guess, you know, for $9 a month...

Leo: Bitmap, yeah. Get one of those Ubiquiti EdgeRouter X's.

Steve: Oh. This little...

Leo: You're going to have to explain - I got mine. But you've got to explain, what do I do with it?

Steve: Okay. Actually, we have a beautiful application note from John Baxter in our Q&A. So we will be covering that.

Leo: Okay.

Steve: However, first errata, first piece of errata - funny you should mention the Ubiquiti. Many people who are in love with theirs and know them well corrected me when I said last week that it supported PPTP and IPSec, but not OpenVPN. It does support OpenVPN, but only from the command line. And I'm still a bit mystified because I got the information that I was repeating from the latest documentation, where under VPNs it lists two, PPTP and IPSec tunnels. However, many people corrected me, and so I'm sure it's correct. In fact, I have in the show notes a link to the step-by-step instructions for establishing an OpenVPN server with TLS encryption using this just beautiful little $49 router. So thank you, everyone. Oh, and someone did report also that it's just Debian Linux in there. So this is also a real little Debian Linux box.

Leo: That makes sense.

Steve: With five interfaces, five physical interfaces, five physical, separate, not just switched interfaces, but logical interfaces that allow you to set up separate LANs, as our application note we get to later will mention.

Second tail-between-the-legs errata is I glibly and incorrectly said last week that, if somebody was using RAID 6, and that failed, that would require three drives to have failed, since RAID 5 allows one to fail, and you continue. RAID 6 allows two to fail, and you continue, meaning that not until three fail are you in trouble. And I said, oh, and that gives SpinRite a great opportunity or great chance of repairing the RAID - and this actually was in response to a testimonial, where someone did this - because it would only have to fix one of the three.

Well, many people who were paying better attention than I was said, uh, wait a minute. It would have to fix the most recently failed one because the other two that had died earlier would have obsolete data. And of course that's correct. So thank you for the correction. And just to clarify, as drives fail, then the RAID goes on without them, and their data is no longer relevant. So two drives fail, your RAID is, like, running now with no redundancy. So when that third drive fails, that's - but maybe it went for a year or six months. I mean, it could have gone for a long time. So that final drive to fail is the one that you would have to use SpinRite to bring back in order to recover the whole RAID. So thank you for paying more attention than I was.

And speaking of paying more attention, I got a correction, believe it or not, that Daleks are not robots.

Leo: What the hell are they? They look like robots.

Steve: I know. They look like stupid robots, frankly, and I hope I haven't offended half of our listeners. I just - I never got into the whole "Dr. Who" thing. I know it's [crosstalk].

Leo: I didn't, either. It started as a kids' show, so...

Steve: Fly around in a phone booth or something.

Leo: Yeah.

Steve: I don't know what's - and now I'm sure that's not really a phone booth. Okay. But it turns out, from someone who knows - oh, and of course you have one on your desk, Leo, a Dalek.

Leo: Yes, it's a 3D-printed Dalek.

Steve: They're supposed to have some weird - they're supposed to have some snorkel thing, too, coming out of the front.

Leo: Yeah, they've got a thing coming - that fell off, yeah.

Steve: Oh, okay, good. Anyway...

Leo: It's retracted.

Steve: Martin tweeted me: "The Daleks are not robots, but malevolent aliens that use a machine to live and travel in."

Leo: That's their spaceship.

Steve: But there's some weird gelatinous thing inside, I guess.

Leo: Yeah, yeah.

Steve: And so it's sort of armor and transportation. So on that one I'm happy to stand corrected. And this is not really errata, but I didn't have anywhere else to put it. A frequent and valued contributor in the newsgroups, Gary Marriott, who is @ramriot in Twitter and in the newsgroups, he just made a comment following up on my discussion of the Facebook abuse protection stuff. He just noted, he said: "Hi, Steve. Facebook Messenger. Because Facebook is the custodian of the remote key that recovers the local decryption key for message logs, this opens them up to being compelled by deception or court order to release that key to expose a person's local message logs, even if the message logs include end-to-end encrypted messages."

And so while that isn't - it doesn't contradict anything I said, it's worth noting that I was pleased by the security model, where when you log your device into Facebook, Facebook provides your key to the device that then decrypts the device's key, which then allows your message log to be decrypted. And so yes, actually, it's one of the things that Gary has really been handy for over on the SQRL side is he's one of the many people who check my work and often finds edge cases which are absolutely worth looking at.

Leo: So it's in the logs. It's encrypted end-to-end. But because it's stored unencrypted in the logs, or no, because they have the key to the logs.

Steve: Yes. So if...

Leo: But that means it's unencrypted and then reencrypted or something; right?

Steve: It's stored encrypted, and the key is destroyed when you log out from Facebook on your device.

Leo: Which no one ever does, of course.

Steve: Yeah. So his point was, if law enforcement - okay. So here's the scenario. Law enforcement obtains someone's smartphone, and they want to know what the secret conversation - this is the point. This is the secret conversation log where you've been using the secret conversation of the new Facebook Messenger. So they obtain someone's smartphone. Now they go to Facebook and compel Facebook to release the key for that user's account. With that key and the phone, they can then decrypt the previous message log. So that's worth paying attention to.

That's, for example, something that Apple has made much more difficult for themselves than Facebook has been able to make, mostly because Facebook just doesn't - isn't positioned in the same privileged "we wrote the software and designed the hardware" position that Apple is. So it's not clear that there's anything better Facebook could do. But as we know, it's hard to completely lock these things down against every scenario.

Talking about SQRL briefly, as a consequence of a note in the SQRL newsgroup, you know, I'm getting near the end of this. And there was a feature that I had designed in the beginning that no longer made any sense. And so before I took it out, I wanted to make sure there was no valid use case for it. And one of the responses reminded me of something that had never - something that I never talked about that I take for granted, but it's worth making clear.

So this person - and this is the end of this person's note - said: "All right. Sounds good. It could get annoying to have to enter the entire," and he wrote, "super long complicated password every time my wife and I switch active SQRL user X times a day." And what I was hit by was that the model of password is very different with SQRL than we're used to. So what I wrote to him, what I wrote back in the newsgroup I'll read here.

I said: "Remember that the classic super long complicated password logic requirement is significantly changed with SQRL. In the traditional non-SQRL model" - that we're all operating under today - "your remote web account is inherently exposed to the entire public Internet. So it's only the secret of your username and password that prevents anyone in the world from obtaining free rein to your online account. So that is what sets the requirement for strong password protection." Meaning that the password has to be strong because the only way we have of authenticating is typically some piece of information like our username or email and a password.

Now, yes, there have been moves to tighten that down, like Google will, if you log in from somewhere that you haven't logged in before, you can require a one-time passcode or that kind of thing. So there are ways that this inherent vulnerability is mitigated, but that's still the fundamental problem. And of course we also have the concern that, if that secret escapes from websites, which is happening now with increasing frequency, then we're in trouble.

"None of that remains true with SQRL. With SQRL, physical access" - this is what I wrote - "to your SQRL cryptographic identity is the first requirement which cannot be bypassed." Physical access to your SQRL cryptographic identity. "Nothing other than a distant derivative of your SQRL identity ever transits the wire. So unlike with passwords, your SQRL identity cannot be obtained from monitoring your login traffic. This is an underappreciated aspect of SQRL. The fact that we are authenticating locally to our encrypted identity by briefly decrypting it, rather than globally to a publicly accessible service, represents a huge difference in threat models."

And I finish, saying: "Consequently, our SQRL passwords only need to prevent the use of our local SQRL identity by someone who can first obtain access to it." So I wanted to make that point, that is, we're - and I've said it often. SQRL becomes a proxy for our identity, able to identify us to websites. And so what we still have is just so that someone walking by doesn't use that or abuse that, we need to authenticate ourselves to SQRL. But what that really also means, then, is that the password can be, now, I don't want to encourage reckless use because maybe somebody would get a hold of your identity. But you know what your own use case is. Does it never leave your house, et cetera. In which case you can make it easier with no decrease in security, understanding that you simply want to prevent anyone from using your SQRL identity which your password decrypts as needed.

Also, speaking of SQRL, I ordered one. There's something called Sticker Mule. I never heard of Sticker Mule before, Leo.

Leo: I haven't either. What is it?

Steve: It's a site where people make stickers.

Leo: Oh.

Steve: And somebody made a two-inch by two-inch beautiful-looking SQRL sticker. I have a link in the show notes. And in fact, if you click on the link in the show notes, it'll take you directly to the SQRL sticker page. Maybe you can just search Sticker Mule for SQRL. Anyway, I ordered one. It's $2.70 in singles. Price goes way down if you order many. But I have on the SQRL pages, long ago, as soon as I got this logo nailed down...

Leo: I like that [crosstalk].

Steve: ...I posted high-resolution line art for the final logo. And that's it. So that would be 2x2. And somebody is saying, hey, you know, spread the word, SQRL. So if anyone wants a SQRL sticker for whatever purpose...

Leo: Sticker Mule.

Steve: Sticker Mule has them.

Leo: Yeah, $2.70.

Steve: Couple of bits of miscellany. Someone sent me a brilliant observation. I was speaking last week about how the new two-part Healthy Sleep Formula component, the key component, niacinamide, has sold out, like in six different online - the major six different online retailers. Someone said search by UPC. And it's like, oh, it's brilliant. So I gave the people in the newsgroup a chance to get any if they needed it first. We gave them a day and then added it the Healthy Sleep Formula web page. So if you're somebody who has not been able to find niacinamide, on the Healthy Sleep Formula page at the top is the UPC for that correct Source Naturals time-release 1,500-milligram niacinamide. And it's all over the place. Not at major suppliers, but you can find it by UPC. So Bob in Santa Barbara, thank you for that tip. That will be really handy.

I already mentioned "Star Trek Beyond" at the top of the show, so I won't - I had that here, a reminder in case I didn't talk about it before. And Eric Ebert tweeted me, he said: "@SGgrc How do you feel about Season 2 of 'Mr. Robot' so far?" And I thought about it for a minute or two, how to best describe my feelings. And I said: "Put it this way. I'm only still watching because Season 1 was so amazing." So...

Leo: There you have it.

Steve: Yeah.

Leo: Yeah.

Steve: Be interesting, when you catch up, Leo, to see what you think. It would be interesting to see what you think. It's, you know, as I said last week, we're spending an awful lot of time rummaging around inside of Elliot's head. And it's like, eh, okay, it's just depressing in there.

Leo: I kind of lost interest early in the first season, to be - not early, but about...

Steve: Oh, you did.

Leo: ...halfway through the first season. I haven't - that's why I haven't caught up. I'm kind of interested, but not that interested.

Steve: Yeah, in that case there's other things to watch.

Leo: You know what you should watch that you might not have seen? It's a Netflix original featuring Winona Ryder called...

Steve: Like her.

Leo: Yeah, who doesn't - "Stranger Things."

Steve: Ooh.

Leo: I think you will like it. It's a little bit X-Files-y kind of creep show.

Steve: Nice.

Leo: But what's really interesting, it's about the '80s. And it takes place in the '80s. And it really is, in a way, a call back to '80s movies. So there's kind of three groups that you follow. There's kids, and they're great. It's kind of like "The Goonies," or maybe a little E.T.-ish. There's Winona and her peer - there's teenagers, and they have their own story. It's all related about the same story, but their own perspective on it. And then there's the adults. Winona is the adult, and there's a great sheriff. And it's got this - it's actually very multilayered and fascinating.

Steve: A good vibe?

Leo: Yeah, and it's well written. It's good. I think you'll enjoy it. Try it. You can binge it. It's on Netflix.

Steve: Nice.

Leo: Yeah, it's "Goonies," "Stand by Me," "E.T." It's got a very '80s - there's '80s music, '80s hair, '80s outfits. And then it's creepy as hell. The only bad thing is the monster is like kind of not - it's like...

Steve: Oh, no, no, no, no, no...

Leo: Never mind. Never mind.

Steve: No spoilers.

Leo: No spoilers.

Steve: No spoilers.

Leo: Just a little tip.

Steve: And this will not come as a spoiler to many people. I got actually a long tweet, looks like, from Ralph Griesenbeck, whose handle is @RandomGravy on Twitter. And he asked me a question. He said: "In a recent Security Now! episode you recommended looking at the SMART screen in SpinRite. However, that only works if it's supported in BIOS. In my experience the systems supporting SMART in BIOS are not that common. Even if SMART data is available, interpretation is a bit of an art as each drive maker has different implementations. Am I missing something?"

And so I replied to Ralph. I said: "Hey, Ralph. Some BIOSes can and do support SMART probing. But SpinRite does its own directly to the hardware, continuously during operation. So it does have access to the drive's SMART data, even when the BIOS is 'not so SMART.' SpinRite also performs some SMART interpretation for the user and succeeds in eliminating some of the drive-to-drive variations, though you're right that differences between drives can be confusing." And then I sent him some links.

I have two links from the SpinRite pages: GRC.com/sr/smart.htm, and then also sr/smart-studymode.htm. And if users haven't seen those, I really commend it to their attention. I broke that SMART monitoring page down with bullets and callouts showing what every little section does and what they mean so that it really clarifies that because there's a lot of information. It's a very information-dense page. And I would argue it's one of the key features of SpinRite. As I mentioned before, if a drive is just idling, doing nothing, the SMART system doesn't really report anything because it reports on struggles, essentially.

Well, SpinRite makes the drive struggle like nothing ever has. And so watching what SpinRite causes the SMART data to do is extremely illuminating. You can see it generating correctable errors and uncorrectable errors. And SpinRite calculates the error rate that is the number of errors per megabyte of data read from the drive to give you a numerical sense for that. It also captures the minimum and the maximum which have occurred within megabyte samples. And they shouldn't vary too much. If the maximum is really high, that means there's an area in your drive where it had to work much harder than the average for that drive. And there's all kinds of other things, all of the health parameters. SpinRite captures the starting health parameter and then shows you any decrease in that health parameter that is created by the work SpinRite is doing with the drive.

So again, I know we've got a lot of fans. If we have any SpinRite fans who didn't really know what was going on there, next time you're running SpinRite, page over through the UI to that SMART page. And maybe compare that or look at the /sr/smart.htm page to see what's going on there because a lot of good info there. And it was funny, when I was doing the work on SpinRite 6, there was some concern about whether drives would balk at being probed while they were busy working. I saw nothing in the spec to indicate that it would be a problem. But nobody else does that. It turns out it's never caused a problem. And believe me, we've got a lot of experience now with SpinRite 6 doing this. So that's become a solid part of SpinRite's core technology moving forward, as well.

Leo: I like your TWiT IPTV T-shirt, Steve Gibson.

Steve: I like it, too. It is sheer.

Leo: Sheer.

Steve: So it's nice on a hot day like this.

Leo: Is it hot down there?

Steve: Oh, boy, yeah. We've had the heat wave through here. It's better, but boy, is it humid. It is so - it's like muggy.

Leo: Well, go for a - you're near the ocean. I didn't even know this. You're near the ocean. Go for a dip after the show.

Steve: Yeah. Hold on a second, I'll be right back.

Leo: Do you ever go for a dip? When's the last time you've been in the ocean? Was it in this decade?

Steve: No. Unfortunately, the ocean's pretty screwed up now.

Leo: Is it? You don't want to go in now, huh?

Steve: You don't want to go in there, no. It's a little scary.

Leo: That's too bad. We're going to be in San Diego in a week. Is it going to be...

Steve: San Diego's got nice beaches.

Leo: Yes. I can't wait. We're going down next Wednesday.

Steve: The problem is there's a lot of sand.

Leo: All right. Just stay inside and finish SQRL. That's all we ask, Steve Gibson.

Steve: That's what I'm doing.

Leo: That's all we want from you. Just keep working. Keep a-working.

Steve: I hear you.

Leo: Question #1, Tod Sage, field support technician. He says the real world may differ a little from our ivory towers in cryptomath: I've been a field service technician since 1988. I've worked for many companies, as well as myself, in a very wide range of disciplines. Now, these settings raise concerns about the level of access I have been given. My biggest concern is when one large company contracts with another, that then contracts with another, and so on, until I'm looking at job postings on Craigslist from some company in India who sends me a packet containing usernames, passwords, and an ID badge stating I am a contractor for Company #2, with instructions not to tell the end customer the full details of who's cutting my paycheck.

For someone replying to a Craigslist solicitation for a one-time job, there is very little to deter any active engagement in cybercrime, and very little in the way of verification even that I am who I say I am. And thanks to the way the hiring chain works, no one involved even knows who I am, and no one really cares. That's somebody else's problem. Having been involved in this firsthand, and having seen this occur often, I have seen no sign of any enforced standards in IT security. Is it any surprise, then, that the many corporations who outsource their IT service needs suffer so many problems and have such poor security track records?

Let's not forget the most famous contractor of all, Edward Snowden, who was a contractor at the NSA. Here, Edward, have access to everything.

Steve: Yeah. You're a smart guy. You've proven yourself.

Leo: What could possible go wrong?

Steve: Yeah.

Leo: And of course we were talking about this yesterday. The Target infiltration went through their HVAC company, that Target gave access to the network to the HVAC contractor. And who knows who they gave it to?

Steve: Right, right. And I just - I loved this story from the real world because we do get caught up in minutiae, like, well, you're never going to be able to factor that prime. But if some guy in India sends you the password, you don't need to.

Leo: Wipro is the company in India that hires a lot of contractors, tech contractors. W-I-P-R-O. They've massive.

Steve: Yeah. And I'm sure that their contract says we're just making our best effort.

Leo: Yes. Do your best.

Steve: And if anything happens, well, you know, it's not our fault.

Leo: Kyle Day tweeted - well, I don't know, maybe this is a DM. Hi, Steve. Love your SN podcast. Highlight of my week. I have a question about corporate spying via inserting a certificate into a Windows user's root certificate store. I understand that IE and Chrome use Windows' built-in certificate store, but Mozilla's Firefox uses its own. If I install Firefox on my work machine and use it for personal browsing, does that mean that it's impossible for my employer to decrypt that traffic because they don't have a certificate to MITM my Firefox traffic? Is that a workaround for corporate spying?

Steve: So a couple things there, sort of some that we've talked about, some that we haven't. So first of all, Kyle's understanding is correct. Firefox maintains its own certificate store. He would have to verify that whatever mechanism the corporation had for getting the certificate into Windows would also not work for Firefox. But it's likely that it wouldn't. That is, the active directory group policy stuff leverages Windows specifically. So using a browser that brings along its own store would work, that is, would prevent anyone like an employer at the border doing a man-in-the-middle.

The problem is that browser may not be able to get out on the Internet at all. That is, it is very likely that anyone setting the system up - and this kind of comes back a little bit to the story we'll be getting to about the EdgeRouter configuration, the application note. That's the word I was trying to come up with, the application note that we have, because a corporation could block HTTPS traffic, requiring, for example, the use of a corporate proxy. And the proxy would absolutely require the presence of a certificate.

So essentially the answer is, it might work. And if it does work, then it should be secure. But if the corporation is really serious about filtering all traffic on their Intranet, then it will not be possible to not go through their proxy, which is also decrypting and inspecting the traffic as it happens. So while the concept is right, it's very common that it would also be blocked, unfortunately.

Leo: Scott Ericsson, Milwaukee, Wisconsin with a SQRL question: Steve, SQRL sounds amazing, but I think there's a problem. How many emails a day do you get with that, that begin like that?

Steve: Yeah.

Leo: I have found a problem. As I understand it, SQRL auto-magically creates a unique identity for each of its users for each website they visit. But what if I want to appear as a different user at the same website? Under the Internet's present insecure email and password scheme, I can use a Gmail alias, or use a secondary email account to create a second independent identity at any site I wish. SQRL would appear to lock us into our SQRL identity for each site. Is that not a problem?

Steve: So, great question, one that we've had before, and one that we have an answer to. There are many ways to solve a couple problems, depending upon what makes the most sense for the user. You could certainly create another SQRL identity. And so, for example, in a household, each of the kids and Mom and Dad would have their own SQRL identities. And nothing prevents you from creating an additional identity for use at a certain site. The problem with that, I mean, and that's absolutely - you can do that.

The problem is there's some overhead that comes with a SQRL identity, like that rescue code that I talked about, where you have to store that somewhere. Essentially, you're doubling up the stuff. When you set up a new device, you would need to import or to export your identities and then import them into the new device and so forth.

There's a better way. Built into the protocol we have a mechanism known as "alt ID," alternate identities. And any time you are authenticating yourself to your client, to SQRL, there is an option button that allows you to do a couple extra things. One of them is to change your SQRL identity, sort of in a sticky fashion, if you wanted to switch to somebody else's identity.

But the other option is to use an alternate identity. The way SQRL creates the identity, remember, is it hashes the domain you're visiting through a keyed HMAC, where the key is your super secret master identity. The alternate ID is simply appended to the end of the website's domain name. So it can be anything you want. It could be the numeral zero. It could be HiMom. It could be anything. And we simply add that string to the end of the domain name, which creates an absolutely separate, unique, non-linkable, non-trackable, it's a completely separate identity for that site. And so it's built into the protocol, it's defined in the spec, and it exists in the client now. So Scott, we've got you covered.

Leo: Oh, I dropped my headphones, hold on. Whatever you're saying, I - okay. Now I'm working all right.

Steve: I ought to also mention, I mentioned last week that there was something else I wanted to talk about, but we ran out of time. We were right up at, like, two-plus hours.

Leo: We've got lots of time today. Go ahead.

Steve: One other thing that I added about a month ago was there was some discussion about a feature in v2 to produce a static secret. But it was so simple to do, and so useful to have, that I said, no, it's going in right now. And the next iteration of the client that everyone got had support for that, and it's in. The idea is, think of things like LastPass. In the LastPass model, they have data they're storing on our behalf, yet they cannot decrypt it. They need something from us in order to decrypt the blob that they're storing. That's a really useful model, the idea that a website could have any amount of anything. I don't mean just a password database, but user data. And they don't want to be able to decrypt it. And if they can't decrypt it, they're not vulnerable to any kind of attack.

Well, right now the SQRL spec that I've discussed had no such provision. We give the site our public key to use to identify us, which does double-duty to not only identify, but also to authenticate because we sign a challenge, and that verifies the signature. But there isn't a secret that we're providing that could be used as the master decryption key for something server-side, until now. It's there.

And so if the site wishes to obtain a static secret from SQRL, in the first exchange it sends what's called a "secret index." And again, that can be any information the site wants - just the numbers, just a numeric zero, or a wave of the hand, it doesn't matter. The SQRL client generates essentially a subsidiary static secret from hashing what the site provides off of the master identity and returns it. So, and that's, like, it's the output of another - it's another 256-bit output of a rather complex hashing process to make sure that there's absolutely no way to go upstream. And it's a little overkill, but it doesn't take up any time, so we go for overkill where we can. And that allows a website to obtain from someone authenticating with SQRL a secret which will never change. Every time they come back, just as their identity is the same, the secret is the same. And the system also handles previous identities and previous secrets in the same fashion. So it's possible for the site to ratchet itself forward.

Basically, we have it all covered. It's one of the things that's taken a while. But you know me, I want to get it right so that I never have to look at it or think about it again, and I can get back to SpinRite 6.1, and SQRL will be able to launch and solve the world's problems. So anyway, it's very cool that it's able, that the site cannot decrypt something it's storing on your behalf. But the SQRL technology, this little addition, simple, took minutes for me to add it, will then allow the website to request a static secret, which then it can use to decrypt data it's storing for us for as long as it needs it. So, very cool.

Leo: From Docop, @docop29.

Steve: Who knows.

Leo: Who knows. Twitter handles, what can you do? Worse than license plates. Steve. I have to figure out what my license plate for my Tesla should be. People do a lot of thing with electrical, you know, WattUp. Amped. Stuff like that.

Steve: Those are probably taken, unfortunately.

Leo: Yeah. All the good ones, I'm sure, especially in California. Anyway: I have a guest WiFi network at work that I use with my iPad for doing mostly work tasks. I have an iPhone that I never connect to the network so my private information is not going through corporate servers. The WiFi password changes every two weeks, so I have to reconnect my iPad. I've noticed recently, though, that my iPhone also connects to the network. It seems that iOS is automatically updating the password across iCloud. Oy gevalt. Given your recent discussions about corporate appliances breaking SSL and being able to access all your traffic, is this going to open up private traffic on my phone to my corporate overlords?

Steve: So, first of all, it is absolutely true that this is one of the things that Apple has decided will be a convenience. So, and for example, for me it's a convenience. As I mentioned, I have about 12 iOS devices. And when I was setting up my new, that Soekris Engineering box running pfSense, I changed my WiFi around. I was setting it up. I came up with a crazy, unhackable password, and I only had to put it into one device. And it was a convenience that all of the other iOS devices that I had suddenly knew how to get onto my new WiFi network. So that was cool. But this guy brings up a very good use case where it's not what you want. That is, where in this case a device that shares an iCloud account is syncing itself through that to obtain information he would like that device not to have.

Now, the good news is there's a switch in the WiFi options of iOS where you can turn off automatically logging into WiFi networks that you recognize. And while having it off the rest of the time might be a little inconvenient because it just won't seamlessly automatically be on WiFi, you would want it off while you were in this corporate setting so that it would get the password, but at least it wouldn't use it without your explicit permission. And I'm thinking, why, who would change a WiFi password every two weeks? First of all, that really argues against it being a big complex password because, I mean, that would just be onerous. And the only reason I could imagine is that they're trying to stay ahead of people giving the WiFi password out. So if people give it to other people, then - or maybe this IT department has run...

Leo: That's probably what it is.

Steve: It's just run wild.

Leo: We're in charge here.

Steve: Exactly.

Leo: We'll just show you. We'll change the password every other week.

Steve: Oh, my god, and everybody hates those guys.

Leo: Well, and we've talked about that. It doesn't necessarily improve security to change passwords.

Steve: No, no, no. I mean, for example, as I said, it's probably - probably have to be weak passwords if you're changing them every two weeks.

Leo: It just encourages people to put post-it notes on their screen.

Steve: Yes.

Leo: But I think with a WiFi password, you nailed it, it's probably people are being - well, he says they have a guest network, though. So I don't know. But it's probably being given out. That's the guest network; right? So people...

Steve: I think, right, "I have a WiFi guest network at work." Yeah.

Leo: Right, yeah, right? So it's the guest password. But we haven't changed our guest password in five years. You know, and I'll tell you what it is: brickguest. It's like, so what? You'd have to be physically here...

Steve: Right.

Leo: ...to use it. I think I'd notice you sitting on the street playing Pokemon Go with my WiFi. Jared is next, @nucleareye. I love the Twitter handles. Security Now! question: What's the big push behind cloud computing and storage, hey? Moving these services offsite makes us more dependent on Internet connectivity and puts our data at a greater security risk, so it seems. I don't like the idea of depending on someone else to access my data, securing it and having access to it. I get it's a nice thing to be able to access from anywhere in the world, but sometimes that isn't necessary. So why put stuff in the cloud? Thanks for SpinRite. I love it. Listen to the podcast every week and love that, too.

Steve: And Jared, I understand your feeling. The good news is, unless you absolutely have to use a service which is cloud-based, no one's making you put anything in the cloud. Drobo is a sponsor now. Drobo is a cloud sitting next to you quietly humming.

Leo: Cloud-free, yeah.

Steve: And with all the redundancy and safety and convenience that you want in just sort of moving something out of the way. However, you're not everyone. And, for example, Jenny is having her laptop backed up by Carbonite, which is a cloud-based service, and thank goodness because she's not - she just wants the problem solved. She doesn't want to get all involved in the details. So from my standpoint, we're living in this rich environment now where there are...

Leo: We have a lot of choices.

Steve: Where there are amazing open, free, low-cost solutions. Cloud is an option. Now, if your backup storage is with you, then there's tremendous advantages to that. But if some catastrophe happens, then you don't have the advantage of physical offsite. So as we've talked about with backup, there's some advantages to physical offsite, and some consequences in terms of performance, just bandwidth access to something physically remote, and security. So it gets more difficult, but it's also very convenient. I just think we're like in this land of riches right now, with mass storage being so inexpensive that everybody's got some.

Leo: You know, I've mentioned this before, but I got my little Linux NUC. It's an Intel NUC. Got it from System 76. It's running, like, stripped down Debian, because I don't need a GUI or anything. It's a server. It's got nothing on it. It's Debian stable, so it's rock solid. And then I put this thing called Sandstorm on top of it. This is my cloud. This is running out of my house. It's HTTPS. They provide that and DynDNS for free.

They have all these apps in the App Store that you can use. These are all cloud apps, and this is all stored on my server in my house. By the way, so much encrypted that you see the grains of the data files. These individual grains are encrypted and stored separately from everything else. So they're completely kind of secure little packets that can easily be transferred. And I'm using BitTorrent Sync to keep this backed up to here. So that's my cloud backup. It's backed up to my work server, or actually work desktop. Actually, I'm backing it up to several different desktops. And even if somebody got any of those backups, those are encrypted and unusable. It's totally, I think, possible to do this.

Steve: Yup.

Leo: And this is free stuff. Sandstorm.io, I'm really impressed by it. I've got a music player, photo sharing. I've got Dropbox-type filesharing. It was an experiment;, but, yeah, you've got the choices. That's the point. But Jared - and you're right. I mean, do you put anything in the cloud?

Steve: Do I?

Leo: Yeah.

Steve: Yeah. I use Amazon S3, and I've got...

Leo: And you encrypt, I'm sure.

Steve: ...all the podcasts and images for various systems are up there. I had the advantage of sort of having my own cloud because I've got the GRC servers in a physical location at Level 3. And so each of my locations backs up to the other. So I sort of have the equivalent.

Leo: Yeah. I mean, I understand most people, most individuals' homes aren't going to have a cloud of their own because, you know. And that's why I got that Ubiquiti EdgeRouter.

Steve: But Drobo is a cloud of your own, essentially.

Leo: Exactly, right, yeah. Torleif Hensvold.

Steve: And this is the one we are going to skip because he's suggesting that why don't I set up a web page rather than have my bit.ly links hijacked.

Leo: Done. Done.

Steve: Good idea. So Torleif, thank you.

Leo: Well done, thank you.

Steve: Done. Link Farm.

Leo: Link Farm. GRC.com/linkfarm. You know that is what Google calls those spammy sites with lots of links on them. But you don't care.

Steve: Oh, really?

Leo: Yeah.

Steve: Didn't know that. Well, mine's going to be lot of links, but it's not spammy.

Leo: It's good links, good links. @scruffydan on the Twitter, Dan Moutal. FYI - oh, this is the answer that you - the errata, the Ubiquiti EdgeRouters...

Steve: Ah, no, the next one is.

Leo: ...do support OpenVPN.

Steve: Oh, right, right, right.

Leo: Need to use the CLI to configure. Speeds not great, 10 to 15Mb. I'm using the POE model EdgeMAX EdgeRouter, which can route at 1Gb. And I tested it, and it works like a charm, he says. Plus, if you really want to geek out, it's just Debian under the hood, and you get full root access. Keep up the good work. That's a different thing than the EdgeRouter X; right?

Steve: Yes. And so I wanted to make sure people knew that they did have a higher performance router that can run at a full gigabit per second. Because I had mentioned that the X will run about half a gig. And so you're getting an incredible lot of functionality for 50 bucks in the X. But they do have a more powerful one that can run faster.

Leo: And now we get to - by the way, I know you're not watching the Democratic National Convention, but I have it on in the background. They're doing the roll call.

And in just a few minutes it will be completed. It's close. Hillary Clinton has 2315 votes, Bernie Sanders 1502. It's like neck and neck. It's a horse race.

Steve: I'm glad they're doing it because the Bernie Sanders voters need to just have that done.

Leo: Well, no, absolutely. You know what, and he got a lot of planks on the platform, as they say. I just - it reminds me of my youth, watching these. "The great state of Montana, home to the cowboy hat."

Steve: With their signs.

Leo: I just love that. I don't know why. It just - to me, that's American democracy in action. Or something. John W. Baxter, Port Ludlow, Washington, provides a terrific real world application example for the $60 EdgeRouter X we've been talking about from Ubiquiti: Steve, it was interesting to see you discuss the Ubiquiti EdgeRouter on Security Now!. The box does indeed have the ability to create the desired isolated network for IOT devices. We're using it similarly to isolate a guest network from staff networks, as follows:

We're using the machine to load balance between two WAN connections. At home I'd prefer to use my DSL in failover mode only, but I haven't convinced the boss. I weight the cable connection at 95% instead. So he's got DSL and cable, and he's using both. Instead of using DSL's failover, they're bonded, but he's doing most of the bandwidth from cable.

Steve: Yeah, and how cool that you can commit two ports to the WAN side so that if either one goes down, the other one just picks up the slack.

Leo: I could do that with that little Ubiquiti EdgeRouter X?

Steve: Yes.

Leo: Wow.

Steve: Yes.

Leo: Very sophisticated.

Steve: Yes.

Leo: Before you get too far into working with the machine, you should be sure to update and install the latest firmware, which is version 1.8.5, as there are several advances in that version of the EdgeOS software. Until recently, TLS connections to the GUI presented an expired self-signed certificate - a little scary and seems to upset current versions of Firefox. I checked just now: Version 1.8.5 uses a self-signed certificate which expires in 2024.

The built-in DHCP server works well. Each LAN gets effectively its own, as can each virtual interface if you create them. We haven't experimented there. The DNS forwarding and management also works well and can be quite powerful. We're using OpenDNS, paid, to gain the filtering of "unfortunate," as they say, IP addresses, and we don't allow the staff LANs to use any other DNS. Just yesterday I configured the company site EdgeRouter - these are bigger routers than the one we're talking about.

Steve: No, this is the little thing.

Leo: Really.

Steve: This is this little cute box.

Leo: Wow.

Steve: Yeah, the little EdgeRouter.

Leo: To permit explicit configuration of the OpenDNS server's IPs, rather than insisting on the 192.168.X.1 forwarding server. For the guest network, we provide Google name servers via DHCP, but allow overrides for any other DNS server the client desires - that's nice, guests, use whatever you want - and we allow the guest network clients to use the OpenDNS servers as an exception to "the guest network can't touch the ER-X" firewall rules.

I do most configuration work in the GUI, using Safari, but I'm getting better with the command line interface. User passwords for the web interface must be installed using the CLI. Perhaps that's been fixed in 1.8.5. OpenSSH works well with key files. We've disabled password login on SSH. I always do that, too.

Steve: Yup.

Leo: I also do some configuration by editing the downloaded config.boot file and installing the result. Terminology: Ubiquiti can't decide whether these things are EdgeMAX or EdgeRouter. Also note that the Ubiquiti community forums have many articles whose details have been obsoleted by newer versions of the software, but they're not dated, so that's hard to detect. Enjoy your explorations of this machine, but first please finish SpinRite 6.1, okay? Come on.

Steve: Okay. So a couple things. First of all, everyone should know I'm not using that. I've solved my problem with the Soekris Engineering hardware and pfSense. And that does everything that this little Ubiquiti router does, and way more. Although that's probably arguable, given that it's running Debian, and you can probably install whatever you want to on it. I mean, it's an amazing piece of hardware for $50 with five physical ports. So anyway, I'm not wasting any time or spending any time. That's why I don't know it better than I do, and I'm just reciting what other people have told me and what the manual says, even when it's wrong. But I wanted to bring it to our listeners' attention because, for 99.9% of the people, it's amazing.

And so just to summarize what John said, as an example of the control that this gives you, he's got two separate networks with completely separate DNS management. The corporate LAN is hardwired to use OpenDNS's paid servers and blocks any attempt not to. So not only do you say here's the DNS you'll need to use, it won't let you make any changes. Whereas he deliberately configured the guest LAN in a more lax fashion. OpenDNS is presented through DHCP. So in that "obtain IP address automatically" mode, you're also getting the OpenDNS DNS servers. But he deliberately said, but if you want to manually configure your own DNS, you can do that, too.

So, I mean, there's so much power in this little box for $50. Anybody who wants to mess around with this kind of next-generation professional-level packet networking - and, as I said, yes, we solved the isolation problem with three dumb routers. But you could also just use one smart one, just this, for $50. And, boy, you'll just - you could play with this thing forever. So John, thanks for sharing the details of the way you set this up.

Leo: One of our chatters, Neo, has mentioned that Dan Gillmor, who's a friend of TWiT and a journalist, has raised issues about the terms of service, the End User License Agreement, for Ubiquiti. They say they can't - I'll tell you what it says, and then I'll say why I don't think it's an issue. But he said they can collect information about you. And that's often the case in terms of service because there's information that you give them when you log into a Ubiquiti account or whatever.

Steve: Right, right.

Leo: But so just to be aware of, some people might find this cause for concern. Read the license agreement.

Steve: Yeah. There's no indication that this thing is sending anything home.

Leo: I don't know how they would collect information from this, anyway.

Steve: No.

Leo: Because it's just a router; right? It's not phoning home.

Steve: No, it's not phoning home.

Leo: So, now, my question is, the problem is I have this Eero. And the Eero really wants to talk to the Eero servers. This is this new WiFi thing. So I don't think I want to put the Ubiquiti in between the Eero and the outside world. But I could put it between - I could put the Eero directly connected to the cable modem, and then off the Eero go to the Ubiquiti and have the server connected to the Ubiquiti, and use the Ubiquiti rules then; right?

Steve: So what's this Eero?

Leo: Eero is a - okay, this is another thing. Someday you want to look at this. This is a new category of WiFi routers that's very interesting, very expensive.

Steve: Is this the mesh system?

Leo: Well, I don't know if it's mesh or not. I think it's mesh, but I'm not sure. Yeah. There's three of them, if you get all three. You distribute them, and it distributes - it does a great job of really boosting my WiFi signal, and there's no dead spots anymore. And they do something else which is you have to establish an Eero account, either with your phone number or your email. And the router is logged into the Eero servers. They update the firmware, like all the time, which I think is a good thing. They also claim that they are doing some tuning. They see what devices you're using, and they're tuning the router based on what the device is. All I can imagine is maybe doing some QoS stuff for video streaming [crosstalk].

Steve: I'll bet your household is keeping it busy.

Leo: There's a lot of stuff on it. It's not nearly as...

Steve: How many light bulbs does he have?

Leo: Yeah. It's not early as configurable as the Asus I use, or this EdgeRouter. I mean, you do it all with an iPhone or Android app, and it's just - you know what I should do is have...

Steve: More turnkey.

Leo: Yes, very turn - it's great for somebody who can afford - I think it's 500 bucks - needs really good WiFi, and doesn't want to geek around with it. No idea what their security model is or anything, although they seem like - the guy comes from Google.

Steve: So two things. The problem with doing a Y connection is that they would each have to have their own public IP. So you would have to have two IPs.

Leo: Ah. So you really do want it to be the first thing on the connection.

Steve: I think so. And the beauty of this router, it might require a little bit of tinkering, but I'm sure you could get it to pass right through so it didn't even know that there was a router.

Leo: That's what I would do, yeah. Pass through to the Eero.

Steve: Yes.

Leo: And let the Eero do its DH - sounds like, from this guy's email, that the Eero can say, oh, you do DHCP. Oh, no, I'm not going to do it for this one. You can really control it that way very granularly. All right. I'll figure it out. There's so much new technology in my house now that I have...

Steve: Arriving daily.

Leo: I have, yeah, kind of have to slowly work my way through it. Tonight's the FreeBSD box. Let's move on. Fred, I'm sorry, Tom Zitzelsberger. Steve, I just - that's his name. Don't laugh at the guy's name. I just listened to SN-569, and I have a question. You said that the new Facebook secure message system was crypto done right. But it seems like it might be simple for Facebook, having received a National Security Letter, to simply add a single character to the secure message to trip the recipient's report feature and have the now-decrypted message returned right to Facebook without notifying either sender or receiver.

Also, since the message needs to go through the Facebook servers, it would be simple for Facebook to append that single byte to all incoming traffic to the person named in the National Security Letter, and the target's own system would return all their incoming private messages, as well. It seems like this new system would fit perfectly with Vladimir Putin's new law about companies providing the FSB - the Russian secret service - with means to decrypt messages. Let me know if I'm mistaken here. I'm a huge fan. Thanks for all you do for us.

Steve: So, yeah. There's a little bit of misconception here that I wanted to clear up, and it's important to understand the way this abuse reporting works. If the recipient's device found that signatures don't match, that is, if as Tom suggests a character were added anywhere, that would bust the hashes and signatures, and the incoming message would be discarded. It would not be echoed back to Facebook. It would just be thrown away. It would be a communications error. The system would assume, oh, there was a transfer error, so this is gibberish. Just don't display it.

The decryption is only - and when I say "only," I mean Tom's suggestion and anything else anyone can come up with, only if the recipient sees the message and is offended by it, it feels that it's in some way abusive and against Facebook's terms of service, then the recipient of the decrypted message can themselves voluntarily choose to bounce it back to Facebook in the clear, along with the various other tokens that we described in detail last week that allows Facebook to validate that that is unchanged from what the sender originally sent because it did go through Facebook encrypted on the way.

And this technology we discussed last week allows Facebook to say, yes, you know, essentially they reencrypt the message that was sent in the clear and are able to verify it's the same thing they originally got from the recipient, who has created this offensive message. But I wanted to make sure that everyone understands, messages will never be seen by anyone but the recipient unless that recipient chooses to break cover and send a message back to Facebook.

Leo: Which could happen, too.

Steve: Yeah.

Leo: Scott Surbrook. I've been watching Security Now! since the 200s, but I don't remember any episodes in which you apply your ability to simplify complex topics with regards to why 256-bit symmetric key encryption like AES is as strong or stronger than 4096-bit public key encryption like RSA, especially since there are approximately 2^1200 primes in a 4096-bit RSA key space. What would be the equivalent - he says RSA, but I think he means AES key size. Thanks.

Steve: So I think we covered this in Episode 199.

Leo: Oh, we've mentioned it. You missed it. You just missed it by one.

Steve: Yeah, he started at Episode 200.

Leo: Just missed it by that much.

Steve: So Scott, and anybody else who's wondering, the idea is these are approximations. And so there is no fixed equivalence, per se. The idea here is that cryptographers use everything they know - and this is the other reason. It's based on assumptions. So cryptographers assume that there is no way of short-circuiting AES. And if that's true, then they know how many keys are available in a 256-bit key space and, with current technology, how long it would take to crack that by brute force, assuming no other solution. Similarly, in the case of RSA, we know how long the public key is and how many primes are available and the rate at which we can try, the rate at which we can guess.

And the reason public keys are generally a lot larger than private keys is that we're trying to, with a public key system, we have a weaker problem that we are trying to prevent the cracker from solving. So the problem itself is not as hard. For example, what, a prime is three and seven, okay, so there's 21. What's the prime factorization of 21? Well, that's easy. The point is that the problem itself is not intractable. It's just difficult. And so public keys are much bigger, typically, that is, asymmetric keys, because the way the asymmetric key systems we've designed so far are difficult, but not sort of the same, like, absolute, there's no way to short-circuit this other than brute-forcing that we do have with symmetric keys. And, famously, Bruce Schneier a long time ago did a chart predicting the rate at which CPU processing power would be increasing and what that meant for minimum key length over time. And there is a chart that shows one column for symmetric and another column for asymmetric.

And so the point is, this is all just seat of the pants, sort of just rule of thumb. So we're in a place today where a 128-bit symmetric key of high quality is good, as is a 2048-bit public key that is properly derived with high-quality entropy. For the foreseeable future, 256 bits, remember, it's only twice as many bits for symmetric, but it's ridiculously more combinations. And 4096-bit, doubling the public key, provides again just a huge amount of protection. And the reason you just don't have really big keys is there is some computational burden.

So, like, every time you use the key, you've got to do this work. And that does go up quickly as the key size increases in the case of public key crypto. So the argument is it's just waste to use a ridiculously large key, when a key that's a lot smaller still provides enough security margin. And this margin is this notion of, well, here's what we know we can do, and we have a margin between that, it's like our safety margin, for how long we want the key to be able to survive an attack. So again, just sort of rule of thumb.

Leo: Number 11 from @grymoire, another Twitter, Bruce Barnett. CryptoDrop sounds great. However, as far as I know, it's not available. Anyone know how or when it'll be available? Free? Commercial?

Steve: I just wanted to make a note, a number of people were excited by my positive review, I mean, surprisingly positive. It's like, okay, I want one of these. The lead author of the whitepaper shot me a note thanking me for the coverage of CryptoDrop last week. And I wrote back and congratulated him and his three co-authors and said, as soon as there's something that our listeners can take action on, like something they can download, please let me know. So as soon as that happens, I will let everyone know.

Leo: Very nice. And last and least, Ryan Young. [Indiscernible] not least. No, no. If you're last, you're least. Ryan Young. Steve...

Steve: At least you're not penultimate.

Leo: Yes. That was this other guy, Bruce.

Steve: That's right.

Leo: Steve, do I need a Level 3 packet switch to segregate my network? Or can a virtual network in DD-WRT accomplish the same thing? I would like to set up a segregated open network with limited bandwidth on my router, but I don't want anything, anything to get into my home network. I love the show and can't wait until Tuesday for a new one every week. Thanks for all the expertise. Well, thank you, Ryan Young.

Steve: Yes, thanks, Ryan. And I want to say again, this is a big problem with virtual LANs, VLANs. It is easy to confuse it with security. It is not. Virtual LANs are useful for administrating a huge, sprawling network and sort of Ethernet because Ethernet doesn't - there's a point at which Ethernet stops scaling well. Because of the nature of packet collisions, if a single Ethernet gets too large, it starts to fail. So virtual LANs can be used to logically segregate Ethernets, but the physical enforcement has to be performed by a VLAN-aware switch, that is, a switch which sees packets tagged for specific virtual LANs and then only sends the packet out of that port of the switch.

The thing to understand is that the VLAN tag is just data. It's nothing magic. It's a slightly extended couple fields in the Ethernet frame which allows this management of the frames. But confusing that with security, that is, isolation, gets you in trouble because if something else can see packets on another VLAN, well, it just ignores the tag. Even if it's not for its VLAN, it's for a different VLAN, it can still see it. So you're gaining nothing for security. And it can spoof packets for the other VLAN simply by changing that tag in the Ethernet frame.

And he mentions DD-WRT, so it's probably in one of the, as I call them, the blue box routers, which are a router coupled to a four, typically, four-port switch. But that's just a switch. That's not four separate interfaces like the Ubiquiti routers that we've been talking about have. It's just a switch. So there is no isolation available between those ports. You need something more.

Leo: You did it again, Steve. Twelve questions, a new world record.

Steve: And right on schedule.

Leo: Right on time. We do Security Now! by GRC's Steve Gibson every Tuesday, about 1:30 Pacific, 4:30 Eastern, 20:30 UTC. If you want to join us live, love it if you do. But if you can't, no problem. No problem because we have on-demand of every show. Steve's got a copy of the audio at his site, GRC.com. He's also got a lovely written transcript so you can read along as you listen. You can watch and listen at our site, TWiT.tv/SN. And you can also subscribe because there's a million ways to get podcasts, whether it's iTunes or Google or Stitcher, I mean, it just goes on and on. And of course all the great TWiT apps on every platform.

When you visit GRC.com, don't forget to pick up a copy of SpinRite, the world's finest hard drive maintenance and recovery utility. He also has lots of free stuff there. SQRL, find out more about that, where they sit in development. If you have questions for Steve, you can ask him here, but you can also ask him on his Twitter handle, @SGgrc. He is open to DMs, to Direct Messages, so you've got several places you can talk to him. GRC.com is the website, @SGgrc on Twitter, and here every Tuesday. Which doesn't seem like enough. But, you know, I don't want to take any more time away from your very important projects. So thank you, Steve.

Steve: My pleasure, my friend. Talk to you soon. Well, talk to you next week.

Leo: Next week.