SECURITY NOW!

Transcript of Episode #546

Router Q&A Follow-up

Description: After catching up with the most interesting security news of the past week, Steve and Leo address three representative questions posed by listeners regarding last week's "Three Dumb Routers" episode.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-546.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-546-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We have a few questions about the router "Y" configuration we talked about last week, plus security news, and the Hack of the Decade. The hack of the - that's what Steve says. It's actually pretty cool. All coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 546, recorded Tuesday, February 9th, 2016: Our Router Q&A Follow-up.

It's time for Security Now!, the show where we cover the latest news in the security and privacy realm with this guy here, Mr. Steven Gibson of the GRC Corporation, of SpinRite fame. Hi, Steve.

Steve Gibson: Hey, Leo. Great to be with you again, as always.

Leo: Yeah.

Steve: For our 11th year of this podcast.

Leo: Geez. So ever since Year 4 or something, we've been doing a Q&A in every other show. Although...

Steve: Yeah.

Leo: ...sometimes the news intervenes.

Steve: Yeah, I like - I think it's been very effective.

Leo: I do, too.

Steve: It engages our audience. It lets them, well, it gives me great feedback into, like, what things had traction, what they care about. So it's allowed me to tune the podcast over the years. In this case, last week's topic, the "Three Dumb Routers" topic, was of overwhelming interest because I think it's clear to everyone, in fact, Brian Krebs has his most recent posting is about this. And it just happened, and I didn't have a chance to look at it, so I've got it bookmarked for next week. But he's talking about the Internet of Things problem. And anyone who's been following security and device security understands that the Internet of Things is going to be a problem.

So last week's idea was to really get down in the weeds more than we have. And several people did note, either by Twitter or email, that, you know, Steve, you've talked about this "Y" configuration several times. It's like, yeah, I know, because it's the one. But we never went into it in this level of detail. And I think, as the questionable security of a growing number of tantalizing things that would be hung off of one branch of the "Y," as that becomes increasingly worrisome but prevalent, because I think Internet of Things is clearly in the process of happening, then the security, the need to somehow isolate those that are almost certainly, I mean, we could say absolutely are going to have problems becomes increasingly important.

As a consequence, last week's episode had a ton of traction, but generated the majority of responses via Twitter and the email bag. So what I've done is - and we sort of ran out of time, and we had a full two-hour episode last week because I wanted to get into this and explain. But there are some subtleties about inter-router routing, or IP routing, that I didn't get to, which raised some questions. So I have three questions that actually happen to all just have come in via Twitter that I selected, which we will put at the end of today's other stuff, the week's news and miscellany and trivia and so forth, in order sort of to keep with the theme; but, again, driven by our listeners.

So, and we had an interesting week. Not crazy. There was of course the famous, what is becoming infamous, I guess, Error 53 that we'll talk about. Comodo has done something, and in my notes I wrote down: "Comodo's Crummy Chromodo Browser." They actually took the Chromium source and created the Chromodo Browser and did something so awful, I mean, that Tavis Ormandy at Google is running around with his hair on fire again, and rightly so, as we'll see. So we'll explain that. And I just feel sorry for people who go to a site and say, "Oh, look, here's a super secure browser from Comodo." And we'll explain why in a minute.

There was an interesting audit of Windows 10. I think I heard you talking about it yesterday on TWiT, how promiscuous Windows 10's communication to the 'Net is. Unfortunately, the Slashdot posting that got all the attention was completely incorrect in what it stated. So it turns out it's not nearly as bad as people who just read the headlines would be led to believe. I have some experience with GWX and a recent Win7 install of mine that I want to share.

Then I ran across the amazing clever hack of the decade. What is this, 2016? Yeah, of the decade. Certainly the funnest, most cool idea in the last six years. So of the 21st Century. And then a little bit of update on what's going on with Zeo and sleep monitoring. And we'll actually then, believe it or not, go back to routers and talk about that some.

Leo: Now, I notice that your show notes don't, as they often do, have an Illustration of the Week. So I'd like to proffer one...

Steve: Oh, please.

Leo: ...for you. This comes from a tweet from Randy Krum, who listens to the show. And you may remember I was saying we need a new diagram for the router, the "Y" router setup that you describe, because your page is outdated. He proposes this one, which you'll all recognize...

Steve: As the flux capacitor.

Leo: ...as the flux capacitor from "Back to the Future." And in fact it is a "Y." He has relabeled. The DYMO labels now say "Guests," "Main Network," "Root Router," and "Internet." But it still does say, and I think it should, "Shield Eyes From Light." So thank you, Randy, for - I really love it that he added the DYMO labels to explain this. So if you want I'll send this along to you. Actually, he included you in the tweet. You probably have it in your pile somewhere.

Steve: Yes, I did see it. Thought it was very cool.

Leo: Isn't that great? Thank you, Randy. Love it. Love it, love it, love it. All right.

Steve: I was just going to comment on exactly that, that there are - it's the nature of the way the market is now that people are going to have Internet-connected soccer balls and milk cartons and, I mean...

Leo: Everything, yeah. And, by the way, there's some value to that. It's not as...

Steve: Some.

Leo: Yeah.

Steve: Exactly. And in this instance the connectivity is clearly part of the core functionality, so it makes sense. In other instances it's like they're trying to come up with something to sell or some way to raise venture capital or something. So it's like, okay. But as you said, there are applications that absolutely make sense.

Leo: Yeah, and I don't really care if my refrigerator's connected to the Internet.

Steve: No.

Leo: In fact, I think that's probably a bad idea.

Steve: Yeah, I think not.

Leo: They'll have to convince me on that one.

Steve: So top of the week is this Error 53. And you guys were discussing it on MacBreak. Cory Doctorow, of course, picked up on it. And I pulled just a couple sentences out of what Cory wrote from his coverage of it from Friday. He said: "According to an Apple spokesperson, Error 53 is an anti-tampering measure designed to protect the integrity of the phone's biometric security system. The lockout is designed to protect users from trusting doctored fingerprint readers that might allow unauthorized access to their phones. But the phones that Apple is remotely killing have not been doctored. They've been fixed. There are many independent service centers for Apple's products where you can get your phone fixed more cheaply than the official rate. Independent service centers also thrive in places where there are no Apple service centers at all."

So I need to back up a little bit and explain that what's happening is that there sort of was what was initially a mysterious Error 53 which the phone was generating after it had been repaired and then stopping being a phone. It essentially bricked itself. And I don't take this as being nefarious or deliberate or anti-competitive. I think this is just good security. And I think we need to - clearly Apple needs to respond to this and needs to figure out how they can achieve their goal of protecting their user while allowing a phone with sensitive content to be securely serviced. I mean, anyone who's messed around with, for example, higher end servers will know that computer cases often have a switch that detects when the case has been opened because physical access is an almost super powerful thing to have for a device. And the same is apparently true here.

Now, I wonder, given everything we've heard about the way the fingerprint reader works and the secure enclave, my sense is that what this may be detecting is an interruption in the connectivity, that is, a however brief disconnection of the fingerprint reader from the secure enclave element. The point being that somebody could insert something in the phone which captures your fingerprints. That is, physical access is incredibly powerful.

And so what I imagine the security-conscious engineers of Apple did, and I endorse it fully, is they said, okay, here's something that we're offering to users with our guarantee that a fingerprint, properly registered, is the only thing that can unlock it, and it's safe. But with that comes an assumption that somebody isn't going to change the design, isn't going to, you know, the perfect classic old-school example is a keystroke monitor inline in the keyboard. If someone has physical access to your computer, they can stick a tiny little chewing gum-size thing between your keyboard and your computer, logging every keystroke you type. We haven't talked about those for years. They've sort of gone out of fashion because now we're all - we're at another level of fancy remote exploits. But that kind of thing could be done to a phone, capturing fingerprints or maybe a little tiny transmitter gets added to it.

And so it may very well be that Apple is quite appropriately detecting when that has happened. So I would argue that this is something that everyone needs to be educated about, and there needs to be some process for handling it. But I don't take the position, and I don't think it's a fair position, to say that this is somehow Apple deliberately stifling competition. In fact, Rene said on MacBreak an hour ago that even Apple's own service has this happening so that, when somebody at Apple service does something wrong, whoops, that can trigger this error.

Leo: But they are in the good position of being able to give you a new phone; whereas a third-party guy probably won't give you a new phone. So, yeah.

Steve: Yeah. So what's the legal position on third-party service? Is it violating, voiding the warranty of the phone? Can Apple say, hey, we're on firm legal grounds here, you shouldn't have gone to, you know...

Leo: I think we'll find out as these class-action lawsuits make their way through the courts. But were I a member of the court, I would interpret it probably as an anticompetitive measure because you could say you void the warranty, but to actually willfully destroy the phone because it's been worked on by a third-party, that would strike me as bad behavior. But, you know, I'm not a lawyer.

Steve: Well, it's not like a lithium-ion battery exploding. You don't have to stand back 10 feet. I mean, you could turn it back on. They just turned it off. They said something has happened which has a potential for violating your security. This is not a phone anymore. This is currently useful for keeping your door open. But they could certainly reverse that.

Leo: I suspect they will. It'll be interesting to see what happens. I mean, as I said on the show, engineers think this way.

Steve: Yes.

Leo: And it's not their job to understand the legal ramifications or to even...

Steve: Or the PR ramifications.

Leo: PR ramifications. They think, hey, it's not secure. You wouldn't want to use it. Whereas somebody might well want to use it and say, I don't care, I want my phone. So, hey, you mind if I interpose another story that we talked about a little bit yesterday...

Steve: Please. Please.

Leo: ...and got your opinion on it. I was doing an interview on Triangulation about privacy. And we were talking about the Amazon Echo. And but it doesn't just apply to the Echo. It applies to so many of the devices in our homes now that are...

Steve: Our TVs are now also listening, yes.

Leo: TVs, the Kinect from Microsoft, that was the first time we heard these fears. Of course our phones. And he had quoted in an article, he had quoted a researcher in privacy and an analyst for Forrester who said that she and her husband were having a conversation. He had been looking on Amazon for suitcases. They were about to travel. They had a conversation about it. She went back to her computer, all of a sudden she was seeing ads for suitcases on her computer. And from that she deduced that the Echo had been listening to their conversation, notified Amazon, and Amazon had modified its recommendations based on that. And for that reason she disconnected the Echo and sent it back to Amazon.

I'm just wondering what your opinion is on these always-on listening devices. My impression is that it's doing pattern matching. All of these devices - none of these devices have enough smarts yet to interpret - maybe a phone does, but most of these other devices don't have the smarts to interpret your speech. So they will send your speech to the server. But they're not sending everything to the server. They wait to be triggered.

So what they are doing is listening and pattern-matching what they hear to their trigger words. That's why it's a limited subset of trigger words. And when they hear that trigger word, then they in effect wake up and say, okay, now, whatever you say after that, we are going to send to Amazon. It was my opinion that it would be obvious if they were sending all the audio back to Amazon. You could look at it with Wireshark. And I also feel like, besides the clear PR hit if people discovered that, they don't want all of everybody's audio sent back to them at all times. Anyway, that's my opinion. What do you think?

Steve: So, okay. There are two ways this could go. I've mentioned to you before that it is not necessary to pause after you say the trigger word. And that was one of the things I first noticed is you can say her name and then immediately start asking the question. But when you use it interactively, there's a delay between you using the trigger word and the little ring at the top lighting up.

Leo: But Google Now works that way, as well. You can say the trigger words and continue to speak.

Steve: Correct.

Leo: And presumably, I would guess, what's happening is that they're now recording and sending it back.

Steve: Well, so what the device is always doing is listening and streaming all received audio internally. And so it certainly takes some processing for them to decide if the phrase, if the key phrase is part of that buffer. And but the idea is that...

Leo: I likened it to a grep.

Steve: Yeah.

Leo: It's kind of doing regular expression matching on this stream of audio, looking for a pattern.

Steve: Well, and you can, from an algorithmic standpoint, you can definitely design an algorithm to understand a large range of speech saying a given word. That is, the word "A-lex-a," those sounds can be set so that the device will respond to anyone saying that, so exactly as you say. Then, when it's been triggered, it begins storing up the audio, no doubt compressing it, and then beams it, almost certainly beams it off to Amazon, where serious big iron decompress the audio, run it through a standard, sort of like a dictation-style system, where the audio is converted to text, text-to-audio conversion, and then run through some sort of neural network to figure out what it was you did.

Now, at the same time, with bandwidth and connectivity and so forth, it's conceivable - well, and power. I would say, I mean, we know that we have text-to-speech available now in relatively small devices. There's nothing that would prevent that device from doing local text-to-speech except training. Training is normally really required for reliable text-to-speech. And there's no training that you have to go through, although there are some phrases that you can, when you set up the device, that you can read to her in order to improve just the general recognition. It's not clear whether that stays local or whether that is sent to Amazon as part of the profile of the users of the machine.

So really it comes down to what you were saying about using Wireshark. Given these allegations, people will presumably look more closely at the traffic coming from that device. If nothing much happens until the trigger word and the question, and then there's a flurry of network activity, then we can assume it's not listening all the time and sending the stuff back. My guess about that particular story is that there was some searching for luggage that was done before the conversation in front of the device. And then some more surfing was done; and, oh, look. From the original searching that was done beforehand, now we're selecting ads and so forth.

Leo: Yeah. I mean, that's what I said. I could consider - I could come up with three or four scenarios that are much more likely and make much more sense than the Echo listening at all times. And by the way, many people have Wiresharked the Echo. Unfortunately, it's SSL, so the traffic is encrypted. But you can tell by the volume of data that it's not sending back audio at all times.

Steve: Correct.

Leo: Nor is the Kinect, nor is Google Now, nor is Okay Siri or any of that stuff. And no one would design it that way.

Steve: No, it would flood. It would create a denial of service attack on the datacenter to have an incredible number of streams of data coming in 24/7.

Leo: And as I pointed out to our guest, I mean, you can - there are plenty of other privacy harms that are genuine. You don't need to make up FUD in order to prove your point. And in fact you damage your case, if you ask me, if you start inventing harms that don't exist because then people question your credibility in general. And so it's unfortunate that this kind of stuff gets around. And, yeah, of course there's nothing to say that there isn't a switch at Amazon headquarters that the NSA could come in and say, "Oh, by the way, we would like to listen to everything that comes out of Leo's house." Of course, I carry around a smartphone which has a microphone, a camera, a GPS device, and it's in my pocket at all times.

Steve: Yeah.

Leo: Frankly, I think the NSA's much more likely to bug that. And I doubt that this so-called "privacy expert" destroyed her smartphones. Anyway. I mean, it'd be easy, for instance, for Amazon to say, oh, we see you have the same address as somebody searching for luggage. We think you might be going somewhere. Or, I mean, there's lots of ways you could think of this as happening without having to make Amazon's Echo [crosstalk].

Steve: Yeah, it may have been booking those airline tickets that was the trigger.

Leo: Yeah, precisely. Precisely. All right. Thank you. I just wanted to get the security expert to weigh in on this.

Steve: Yeah, I completely agree with you on all points, except that technically you could do local text-to-speech, although...

Leo: Oh, of course you could. But if you look at what's in the Echo, it's not nearly powerful enough to do that on a regular basis.

Steve: So Comodo, who bounces into the doghouse from time to time, they've done something called Chromodo, which is their Chrome version. They took the Chromium source and basically said, oh, we're going to have our own web browser, yay. And so it's called the Chromodo Private Internet Browser. And on their web page it says: "Fast and versatile Internet browser based on Chromium, with highest levels of speed, security, and privacy."

Now, okay. One wonders why would anyone use some random browser when you get the real browser from Google and Chrome, and it's free. But I'm sure, again, this is for the people who our listeners provide support for who are using them for tech support. Turns out this thing is not only not the highest level of speed, security, and privacy, but it's arguably the worst. Google's Tavis Ormandy talked to Comodo. They didn't fix the problem. The 90 days expired. And so, as happens, this went public.

Tavis wrote in his posting: "When you install Comodo Internet Security" - okay, now, this is the other thing, is they're advertising their browser, but it does this to you. "When you install Comodo Internet Security, by default a new browser called Chromodo is installed and set as the default browser. Additionally, all shortcuts are replaced with Chromodo links, and all settings, cookies, and so forth are imported from Chrome. They also hijack" - and no doubt if you were using Firefox or IE, I'm sure they have a generalized, what browser were you using? Now we're going to suck all that in and import it and take over.

"They also hijack your DNS settings, among other shady practices," writes Tavis. "Chromodo is described as 'highest levels of speed, security, and privacy,' but actually disables all web security." He writes: "Let me repeat that. They disable the same-origin policy."

Leo: Is that the third-party cookie thing? What is that?

Steve: Oh, Leo, it's worse. Okay. Which is the perfect question to lead me in. Absolutely, without exception, the only way using browsing on the Internet today is secure is every single, except Chromodo's, rigid enforcement of same-origin policy. And to use a quick analogy, you can think of it as utterly rigid domain stovepiping. What same-origin policy enforces with absolute rigidity is that script running from a given origin is only able to manipulate or access resources from the same origin.

So, and this has become important. Once upon a time, when you used GRC's web pages, for a long time there were no - actually even now. I don't think I have any third-party stuff on my - oh, I might. No. There's the Google search box I think may be the only third-party thing. But the Wayback Machine websites, all of your content came from one place. It came from the site you were visiting. They sent you the images. They send you the media and the pages and so forth.

Of course, that's all gone, as we all know. Anyone who's used Ghostery or uBlock Origin or, I mean, anything, you go to a site now, and 50 different sites are pouring stuff into this page because the page you go to refers to content on all those other origins, all those other domains. So the only way that is secure is if every single item that gets loaded is stovepiped to its own domain. It can't touch, it can't do anything to any other domain, the domain's cookies, or the domain's assets. Chromodo doesn't enforce that. It's probably the only browser in the world that doesn't. It's just - it's a catastrophe.

And these jokers did the same thing with Tavis that we've seen before. He gave them a proof of concept that was one particular example that he just threw together in a few lines of JavaScript, demonstrating how he could, well, for example, you could obtain all of the session cookies for all of the other domains that a user is logged into and send them to some third-party malicious site. I mean, it's just - it's unbelievably awful. And so they disabled one of the function calls that Tavis happened to be using, some exec function, even though, very much like a similar issue we covered a couple weeks ago, you just do it a different way. So they didn't solve the problem at all. They just killed his particular proof of concept.

So for what it's worth, Comodo Internet security brings along with it and by default takes over your system's current browser with the Chromodo browser that does not enforce same-origin policy, which is, I mean, there's just no words. I mean, you're just giving your system to the bad guys. It's like it's worse than any, oh, maybe this could be a remote code execution. No. Any script from any other site that runs on that page could access all the cookies in your browser from all the other sites you visit. And since cookies are the way we maintain login state, they could then impersonate you on all the sites that you visit, or do anything they want to. It's just this is like - it's just unbelievable. And they've blown him off. They said, eh, you know, thanks, but here, we broke your proof of concept. And he's like, well, okay. Oh, it's just unbelievable.

And this is, frankly, this is one of the problems with doing something like Chromium.

Many people have said, "Hey, Steve, why don't you open source your freeware?" It's like, because I don't want people to take it and make it malicious; or say, hey, look, here's a DNS benchmark, and embed malware in it; or just do a bad job. I'm sure these people aren't malicious. I have no idea how this could have possibly happened. But they don't seem to have gotten a clue when someone with a strong reputation said this is really bad, and they just ignored him. So, wow. Yikes.

Leo: And it's a crappy name, too.

Steve: Yes, it is.

Leo: Chromodo.

Steve: Chromodo.

Leo: Sounds bad, just the name.

Steve: It's really - yeah. Oh, wow.

Leo: This is now - we've got to make this clear, by the way. There's some confusion because there's also Komodo with a "K," which is an Internet firewall company. That's not the Comodo we're talking about here; right? This is the marketing company that put the stuff on the Lenovo laptops.

Steve: Correct, yes.

Leo: Right. That's not the same.

Steve: They sell security certificates also, although I'm not buying any.

Leo: Oh. Maybe it is the same. No, yeah. No, I think there's Komodo with a "K." Isn't that a fire - the Komodo firewall, like a Komodo dragon? I might be mistaken.

Steve: Yeah. I know that there - I'm sure there is a Komodo with a "K." I just - I'm staying away from anything that sounds like them.

Leo: This is, no, it is the same, Comodo with a "C." They do a firewall, as well. Well, I guess I won't use that.

Steve: Please, please, please, please, please, please. So Google is increasing the breadth of their deceptive website blocking, which I think is a good thing. It's always sort of controversial when something that sells itself as a search engine, which is just supposed to be indexing what's there, gets into the editorial business, starts saying, eh, we don't like this, and we don't like that. I mean, that's - it's dicey. But the reality is it's a good thing to help users be protected from clicking on links that they are getting, after all, from Google through Google's facility. And so if they click on a link that they got from Google - oh, and by the way, on that page are a lot of ads. That's generating revenue, so it's a revenue-generating enterprise that works because Google's providing you with links. Does Google have no responsibility for what happens to you if you click those links?

So what's happened is they've started blocking websites that use what Google calls "deceptive content," or ads which would cause you to do things you wouldn't normally do, like fake download buttons that are confusing and appear right next to the real download button, or pop-ups informing you that you need to download an updated media player to view the site's content. So it's hard to argue that that's not a net benefit for the Internet's users. And we've seen this before. This is extending the definition, though, from malicious content to deceptive content.

Nobody, I think, would argue about being blocked from clicking on a link that is taking you to a site with known malware that will infect your computer. It's like, thank you very much. This is beginning to blur the line. But still, so many people who don't know better click the link. Up pops the window that says, oh, you need to update your media player.

Leo: Drives me crazy.

Steve: Yeah.

Leo: I've encountered it a lot lately because I've been playing a lot of Minecraft, and Minecraft mods. And when you are downloading the mods, these guys who make these, you know, they don't have any bandwidth or download support. So they use these kind of sketchy downloading sites that are loaded with these buttons. And you have to really carefully parse the page to figure out which of the download buttons is the one you want. It's horrible.

Steve: And they're trying to also get you to use their installer. Well, it's like, use our installer. It's like...

Leo: uBlock Origin blocks SourceForge, which, I mean, SourceForge used to be the place to get open source software.

Steve: I know.

Leo: But now there's so much cruft on those pages.

Steve: Remember that CNET used to be good, and Download.com?

Leo: Downloads.com is terrible, yeah.

Steve: You can't use them now. We've just lost them.

Leo: Yeah. I always tell people, get the software from the person who made the software, if you can. Sometimes, though, these guys are smalltime guys with no bandwidth, and so then you have to go to these downloaders.

Steve: Yeah, you do want to go to the source. And of course this is where the advice of never download something you didn't yourself go looking for, don't, I mean, and again, our audience knows. And unfortunately the rest of the world is going to get caught out by this. It's like, oh, I mean, they don't know how computers work. They're just like, you know, I famously have a friend who thinks Google is the Internet. When I want to talk to somebody about a URL, she says, "No, just use the Google. That's the Internet." Okay, Judy, fine.

So they're calling it Safe Browsing tech, and it does put up, it's a full interstitial page, a big, red, scary, you know, it's the red - instead of the Blue Screen of Death, this is the Red Page of Warning from Google that just says this is a site that is not good.

Now, the good news is this puts some incentive on sites that have not completely gone over to the dark side to come back because, if Google blocks you this way, your click traffic just collapsed. So if Google decides that they don't like the way your site is behaving, you're not getting any links from Google anymore, and that's going to put a serious dent in your click traffic. So it is a nice way of Google using their market power in a way that's arguably helping users to also fix the sites that are behind this big red screen that prevents you from going there anymore. Those sites will have to change their behavior if they want Google to allow their users to come through. And I don't know how they couldn't.

And in one last piece of little Google news, this sort of follows on what Amazon has said. Starting September 1st, Chrome will start blocking all Flash content that it decides is not "central to the web page," in quotes, central to the web page. So Flash content such as ads or auto-playing videos on non-video websites will be automatically paused by default. So it's the autorun-ness of them that, starting September 1st, Chrome is just not going to let them run. You'll still be able to click on them to play them, if you wish. And so embedded video players on sites like YouTube, well, especially YouTube, but also Vimeo, will still work. But Chrome is, again, going to take a look at the page that you're asking it to display and say, eh, we're just not going to let those Flash things run.

So this is, again, more pressure against Flash and moving that same content over to HTML. It's still a majority of ads that are animated are being served in Flash, just for reasons of inertia. And, boy, if this podcast teaches us any lessons about inertia, just say "IPv6." Like, yeah, inertia.

Leo: Inertia.

Steve: Or RC4 or SHA-1 or any of the things that we talk about that people just do not want to give up.

Okay. So the posting on Slashdot by a person whose handle I cannot read, I wouldn't read even if this was on a cable TV channel.

Leo: Oh, dear.

Steve: Yeah. Click on that first Slashdot link, Leo, and you'll see why I'm not giving you his name.

Leo: Yeah. People use weird handles on Slashdot.

Steve: Oh, boy.

Leo: Yeah.

Steve: Anyway, he writes, and this got picked up widely because it was so inflammatory: "Even with telemetry disabled..."

Leo: Huh. I just read his handle.

Steve: Yeah. "Even with telemetry disabled..."

Leo: I think it's expressing his disappointment with this news.

Steve: Or [indiscernible].

Leo: Yeah.

Steve: "Windows 10 talks to dozens of Microsoft servers." And so on Slashdot he posts: "Curious about the various telemetry and personal information being collected by Windows 10, one user installed Windows 10 Enterprise and disabled all of the telemetry and reporting options. Then he configured his router to log all the connections that happened anyway. Even after opting out wherever possible, his firewall captured Windows making around 4,000 connection attempts to 93 different IP addresses during an eight-hour period. With most of those IPs controlled by Microsoft, even the Enterprise version of Windows 10 is checking in with Redmond when you tell it not to, and it's doing so frequently," he writes. Except none of that is true.

Leo: Oh. Well, I'm looking at the log. What's going on?

Steve: This guy, whose name I can share, CheesusCrust - although that's bad enough, but still not as bad as the Slashdot poster. So let me tell you what he explains. And I read it very carefully several times. So CheesusCrust writes: "Like many of you" - oh, and this is on Voat, V-O-A-T, dot co. "Like many of you, I am concerned about the telemetry, spying, and other surveillance features, known or unknown, of Windows 10. It has concerned me enough to push me to Linux Mint as my main operating system. Even so, I wanted to better understand Windows 10, but Internet search results for a decent Windows 10 traffic analysis leave a lot to be desired.

As such, I decided to do my own investigating of what, exactly, Windows 10 is doing traffic-wise, and post the results. For this analysis, I wanted to simply analyze the network traffic of Windows 10 on a clean install, and just let it sit and run without using it." In other words, he disabled nothing. So...

Leo: Oh, this is the default express settings.

Steve: Which we know causes you to need a larger bandwidth connection to Cox or your ISP.

Leo: Well, no, no, no, wait a minute. He says, "I've chosen the customized installation option where I disabled three pages of tracking options."

Steve: Where, where, where?

Leo: Right here.

Steve: Oh. How did I miss that?

Leo: So, yeah, I had read this, as well, and did note that. Now, and by the way, this is Windows 10 Enterprise.

Steve: Oh.

Leo: This is, by the way, not the Windows 10 you're getting if you do the free upgrade. This is a different version.

Steve: You're exactly right. He says, you're right, "I have chosen," in the third point. I just missed it because it wasn't a bullet point. He says: "I installed VirtualBox on a Linux Mint laptop and installed Windows 10 EnterprisePNG on VirtualBox. I have chosen the customized installation option where I disabled" - well, thank you, Leo. It completely changes the nature of this report.

Leo: Well, but also the fact that it's Enterprise changes it, too, because this is not the version most people are going to use. I'd love to see this - but you know what I got out of this? How useful it is to run DD-WRT because you can turn on this kind of analysis at the router. You don't even need to use Wireshark.

Steve: Yes.

Leo: You can monitor all these connections. I mean, Wireshark lets you see what's going on.

Steve: Yes.

Leo: But you can at least monitor these connections. And 1,619 connections to one particular Microsoft address in one eight-hour period.

Steve: Right. So, yeah, so he posts the eight-hour network traffic to 5,508 connection attempts and then breaks it all down. He runs it into a MySQL database and then resolves the IPs, does reverse lookups, builds a whole table. I've got the link in the show notes for anyone who's interested. And it is comprehensive. And so I guess the presumption is an Enterprise version would be less snoopy, being corporate, less snoopy than a home version.

Leo: Maybe. Maybe not.

Steve: But again, who knows.

Leo: Who knows?

Steve: Yeah.

Leo: One of the things that we know happens, if you turn off all the tracking, you don't use Cortana, you don't use Bing, you turn all that stuff off, and that you can do in those customized settings, you still have Windows telemetry, and that's not just on Windows 10, that's 7 and 8, that sends information about how you use Windows back to Microsoft for analysis. And I would presume a lot of this, you know, what is interesting is a lot of these addresses are Akamai. They're connecting to a CDN.

Steve: Yes.

Leo: So that's not typically where you would send data. That's where you would get data; right?

Steve: Yeah, generally that's the case.

Leo: So that's not phoning home.

Steve: So it may have been installing a series of patches and updates over time.

Leo: Drivers, patches, who knows, yeah.

Steve: It may settle down.

Leo: Yeah. I think that I'd like to see more of this. And frankly, I'd like to see some Wireshark stuff. And some of this is port 80, which you could just throw out; right? I mean, that's just browser stuff.

Steve: Well, actually the port 80 stuff means that you'd probably, if it didn't bring up their own encryption tunnel, it would be in plaintext, and you could get some sense for what it was that was going back and forth. Although a lot of it is also 443, so that's over TLS.

Leo: The most connections is over port 3544. That's probably the customer experience program, the telemetry.

Steve: And that, yeah, those are UDP packets over that port.

Leo: Yes, exactly, yeah. It's interesting. I mean, it certainly [crosstalk] data point.

Steve: Yes, it is very chatty. Yup, another - and, again, I'm not - everyone knows how I feel. I will never use it. So I'm not afraid of it. People are welcome to use it, if they like, and knowing what's going on.

After the news dropped a couple weeks ago that Microsoft was actively planning to militate against future chipsets running anything other than Windows 10, and essentially not supporting, not taking the effort to support retired versions of the OS, and given that I'm committed to never going beyond Windows 7, I have purchased a state-of-the-art Haswell chip and all of the components and begun to build my next machine, which is intended to be my last Windows machine. And I think it probably will. You know, the one I'm using has lasted me a decade. I think my next one that'll run Windows 7, it'll last me a decade. That'll take me into my 70s. And at that point you'll be the only one I ever talk to, Leo.

Leo: And you'll be using a PDP-8 most of the time anyway.

Steve: Oh, I'm going to switch to a Mac. The moment I'm no longer a Windows or Intel code developer, once...

Leo: Why not BSD? Why not...

Steve: I think I probably will. I mean, I am a FreeBSD person.

Leo: I think you would like - I think you would really like it, to be honest with you.

Steve: Yeah, [crosstalk].

Leo: Or Linux. I love Mint Linux. And it's, frankly, it's absolutely as usable now as Windows 7 was. But I know you need - because you're coding for Windows environments, you can't...

Steve: I am. And I have a lot of server-side code. And so running Windows 7 synchronizes me with Server 2008 R2, which is what GRC is now on, and where I also intend to stay. The reason I bring this up is that I have a complete experience now with Microsoft's patches to prevent the Windows 10 upgrades of Windows 7. And they work beautifully. We talked about it a couple weeks ago. They're not for the timid. You need to install one or two patches which are not part of the mainstream patches. Microsoft's not giving them to you, either recommended or optional. You've got to go get them. And we covered it a couple weeks ago. There's one for Windows 7, and a different one for Windows 8.1. When you install those it adds features to the group policy and registry that allow you to completely shut off permanently all of the GWX, all of the Get Windows 10 nonsense.

So I set up a brand new Windows, I installed Windows 7 Ultimate on this box that I've built. And first thing I did was, before I installed any updates, it was SP1, Windows 7 Ultimate SP1, first thing I did was I installed that one update for Windows 7 and then flipped a switch in group policy saying "Turn off updates for Windows 10." And I added a couple keys to the registry. And a value had already appeared from the patch that was set to one to prevent it.

I added one more key, a GWX key. Then I went hog wild, installed hundreds of updates since Service Pack 1, bringing the machine completely current, and no sign at all, ever. And it's been up for a week of any attempts to mess with Windows 10. So I wanted to let people who feel as I do know that, if you do that, you don't need a third-party app, you don't need anything else, just that, and it really does shut this down.

In the ongoing 2016 battle, 2017 probably, about what we're going to do about encryption, I wanted just to note that the conversation has escalated from various state legislators to one of our major national senators, John McCain, who gets a lot of press and attention. He did a guest editorial, which I'm sure he didn't write; I'm sure he staffed it out. This was a guest editorial which appeared last Friday, February 5th, in Bloomberg View.

And there was nothing unpredictable about it. It did have sort of a predictable anti-Obama political component and spin, chastising our currently presidential administration for the stance they took last summer of deciding to drop the whole encryption issue for the time being, or maybe having conversations with Silicon Valley, but not going any further. And of course it completely ignores the fact that encryption is already an available third-party add-on, which says it doesn't matter whether your platform's encryption is strong end-to-end encryption, which is now the jargon everyone is picking up. The politicians are saying, oh, it's the end-to-end encryption that's a problem.

So the argument in McCain's posting or guest editorial in Bloomberg View is end-to-end encryption needs to be, if not outlawed, it needs to be altered so that the government is able to get in. And this reminded me, the one thing I wanted to remind our listeners and ourselves is that metadata is still floating around. Metadata is not content. And whereas strong end-to-end encryption protects the content, one of the things we have, we are continually seeing, is that the Internet is not designed to protect the fact of communications. We now have the ability to completely and robustly make them private, that is, the content of them private, but not the fact of them.

And we spend a lot of time talking about Tor and the amazing difficulty that exists in keeping the fact of communications secret. It's just it's incredibly difficult to do that because that's not something the Internet is good at doing. It wasn't designed to do that. We're able to layer encryption onto existing connection stream technology, but not somehow obscure the connection stream. So metadata is still there, still available, and, I would argue, a very powerful surveillance mechanism that isn't thwarted at all by securing what goes through the connection.

Leo: Back to Steven Gibson and...

Steve: Okay.

Leo: Yes.

Steve: This is just - this is just so great.

Leo: Hack of the week.

Steve: Hack of the decade.

Leo: Of the decade.

Steve: Oh, Leo, you're going to love this one, too. So this is from the Check Point blog, where they announce an eBay exploit which is quite severe, which eBay says, eh, we don't think so. We're not going to bother with this. And, boy, it's so clever, and eBay is wrong. So Check Point writes: "Check Point has discovered a severe vulnerability in eBay's online sales platform. This vulnerability allows attackers to bypass eBay's code validation and control the vulnerable code remotely to execute malicious JavaScript on targeted eBay users."

Leo: Oy.

Steve: Yeah, really bad. If this fault - but wait till you hear how. That's the fun part. "If this fault is left unpatched" - and Check Point already informed eBay, and eBay said no. So Check Point said, okay, this'll make a great blog posting.

Leo: Yeah.

Steve: "If this flaw is left unpatched, eBay's customers will continue to be exposed to potential phishing attacks and data theft. An attacker could target eBay's users by sending them a legitimate page that contains malicious code. Customers can be tricked into opening the page, and the code will then be executed by the user's browser or mobile app, leading to multiple ominous scenarios that range from phishing to binary download. After the flaw was discovered, Check Point disclosed details of the vulnerability to eBay on the 15th of December 2015. However, a month later, on January 16th, 2016, eBay stated that they have no plans to fix the vulnerability. The exploit demo is still live."

Okay. Now, remember that - and we've talked about this in many contexts through the years. It is crucial that any site which allows someone to post content which will be shown by that site, to be sanitized. The famous example is [Bobby] Drop Tables, where "drop table" is an SQL command that you don't want your SQL Server backend to interpret when it is processing a web page. But the broader example is you don't want users to be able to post their own content which your site will display because it could be malicious.

Leo: It's the xkcd comic of little - all you have to do is Google "Little [Bobby] Drop Tables," and you'll...

Steve: Yup.

Leo: You'll find it right away.

Steve: So, the vulnerability. "Check Point security researcher Roman Zaikin" - and I'm reading still from their blog, then I'll go into details - "recently discovered a vulnerability that allows attackers to execute malicious code on eBay users' devices" - and I ought to mention, and our listeners will quickly realize, maybe not only eBay, this thing has broader consequences - "using a nonstandard technique called" - and this is not subject for work or whatever that NSFW - "JSF**k" is the way they're naming it. "This vulnerability could allow cybercriminals to use eBay as a phishing and malware distribution platform. To exploit this vulnerability, all an attacker needs to do is create an online eBay store, which is to say their own pages, eBay pages. In his store details, the attacker posts a maliciously crafted item description.

"eBay prevents users from including scripts or iFrame tags by filtering out those HTML tags. However, by using this JSF**k, the attacker is able to create code that will load additional JavaScript from his server. This allows the attacker to insert remotely sourced and controllable JavaScript that he can adjust to, for example, create multiple payloads for different user agents. eBay performs simple verification, but only strips alphanumeric characters from inside script tags. That's the key. eBay strips alphanumeric characters from inside script tags. The JSF**k technique allows the attackers to get around this protection by using a very limited and reduced number of characters."

Now, the technique was invented by a guy named Martin Kleppe, K-L-E-P-P-E. And it is so wonderful. It uses nuances of the way JavaScript interprets empty sets and Boolean values and typecasting. It only uses six characters: open bracket, close bracket, open parens, closed parens, exclamation point, and plus. So the open and closed square brackets, that's JavaScript's array element and object properties and things like numbers of elements in strings, operators. The open and closed parens are used to call functions and avoid parsing errors, so to create explicit groupings. Plus is used to append strings and also to sum elements and cast elements as numbers. So if you put a plus in front of something that's a string, it assumes, oh, they want me to convert that from a string value like a one, two, three string into the numeric value 123. And an exclamation point casts anything following it as a Boolean.

So with that understanding, you can use a ton of just those six characters to essentially do anything you want. For example, to get a numeric zero you put a plus sign in front of an open bracket closed bracket: +[]. So the open bracket closed bracket is an empty list. And when you cast that to a integer, JavaScript gives you the number of elements in the list, which because it's empty would be zero. Similarly, if you do an exclamation point and then open bracket closed bracket, that yields a JavaScript value of false: ![]. Two exclamation points inverts the first exclamation point and gives you a value of true: !![]. If you put a plus sign in front of a true value, JavaScript casts that to a one: +!![].

So now we know how to get a zero with a plus open bracket closed bracket. We know how to get a one with a plus, two exclamation points, open bracket closed bracket. We can cast that to a string by following it with a plus open bracket closed bracket, and so on.

In the example, there's like this crazy string of plus open paren open paren plus exclamation exclamation open bracket closed bracket plus open bracket closed bracket closed parens plus open parens exclamation point plus open bracket closed bracket plus exclamation point exclamation point open bracket closed bracket closed parens plus open parens - and I'm about halfway through: +((+!![]+[])+(!+[]+!![])+(!+[]+!![]+!![]+[])).

Leo: Okay.

Steve: You continue with that for another about that far, and you get the integer 123. So despite the fact that alphanumerics were excluded by eBay, this technique...

Leo: Wow.

Steve: ...taken to a crazy extreme - and there's an example on the link in the show notes, if anyone's interested. And Leo, you can show it if you...

Leo: I'll show it, yeah.

Steve: The link is, yeah, up on the previous page, Amazing Hack of the Decade. You click there, you can see, get a sense for it. It is an incredibly long - and, again, someone had to really want to do this, but you can. And this allows anyone - that's what it looks like. It's just crazy.

Leo: It's a Pastebin of a very long string with this bracket bracket bracket bracket.

Steve: Yeah. All it is is a very clever, patiently constructed concatenation of combinations of six characters that causes JavaScript to turn this into a regular JavaScript code that does anything an attacker wants. And eBay says, eh, yeah, we don't think this is a problem.

Leo: What?

Steve: They'll be changing their tune.

Leo: What?

Steve: The other problem is there are many sites which don't use a whitelisting approach, they use a blacklisting approach, where they'll say, oh, wait, is this a script tag? Okay, don't allow it. Is it a this tag or a that tag? And so they disallow tags they know are dangerous, but they default to allowing other things. This technique will probably see far wider use very quickly. It's very clever, and it's very dangerous.

Leo: Wow. And it's really...

Steve: Wow is right.

Leo: ...really interesting, yeah.

Steve: Isn't it cool?

Leo: Yeah.

Steve: And they go on to show how they can get a subset of letters, and that those letters then, they concatenate those to create function names, and then those function names with appropriate arguments allow them to leverage still further in order to - until they get to the point they can do anything they want to.

Leo: Basically using arrays and parentheses for calling functions, plus append strings and the negative, the not, the bang to cast elements to Booleans. You can generate anything.

Steve: Yes.

Leo: Some sort of obfuscated fashion.

Steve: Very, very clever.

Leo: That's really smart. It's a good code, by the way.

Steve: So I wanted to quickly cover some errata. Last week I misspoke, and our sharp-eared listeners said, Steve, you said "Ethernet packet." We know you meant "Ethernet frames."

Leo: Doh.

Steve: It's like, yes, I did mean frame. So thank you.

Leo: What is the difference between a frame and a packet?

Steve: Well, a frame is the proper term for the equivalent of an IP packet. So an Ethernet...

Leo: Right. Oh, because it's not IP and it's not routed.

Steve: Correct.

Leo: Yeah.

Steve: And a packet would be contained in an Ethernet...

Leo: In a frame.

Steve: In an Ethernet object. And the name of Ethernet objects are frames...

Leo: Got it.

Steve: ...rather than packets. So I stand corrected, and I wanted to, in case I confused anybody else, let them know.

Now, Leo, you'll remember that I talked about an auction for the very cool Zeo Sleep Manager Pro...

Leo: How many did we sell?

Steve: ...at the beginning of last week's podcast. Excuse me?

Leo: How many did you sell?

Steve: Okay. As you'll remember, by the end of the podcast almost a hundred had sold from that auction. That is, that auction was sold out.

Leo: In other words, by the time you downloaded the podcast, there weren't any available.

Steve: And that was a concern because I wanted any of our listeners who wanted them to be able to get them. The auction people, URT Outlet, posted 355 more.

Leo: Wow.

Steve: They all sold out.

Leo: Wow.

Steve: They then posted 460 more. They all sold out by 1:30 p.m. on Friday.

Leo: See, we could have saved Zeo if they'd just advertised on this show.

Steve: Oh, I know.

Leo: Wow.

Steve: That was a total of 915 sold from the beginning of last Tuesday. And I haven't looked, but they just yesterday, and I'm looking for - because I closed Firefox.

Leo: They must have just bought tens of thousands. I mean, do we know how many they got in the bankruptcy auction or whatever it is they...

Steve: I don't have a number. Let's see. Last time I looked they were at 236 of a new auction. And I just refreshed the page. They're now at 253. So here's what I wanted to tell our listeners. Anyone who was excited but happened to go when they were between auctions, they're back. And they have plenty of them. I have updated the GRC.com/zeo page with the link to the auction. They hit some sort of limit that eBay imposes, so they had to switch to one of their alternative identities. But they're the same people. And I just wanted to share this nice note that gives you a sense for them and how impressed I have become.

And I got this a couple days ago: "Good morning, Steve." I think this was like maybe even Tuesday, before the most recent things happened. And this is Dawn, who is their asset sales supervisor. She said: "Yes, this is an awesome phenomenon. We only have matched sets right now. We do not have any more single headbands." So that's worth noting. For a while they had the whole kit for 40 bucks, brand new surplus, and then individual headbands. They sold out of those. And since the pod doesn't work without the headband, they can't sell any headbands alone, at least for now.

And then Dawn said: "I did post another 460 this morning. That is an accurate number for listing, as we are counting daily so we do not oversell this product. I think there are about another 300-plus after the 460, but I need them counted to be sure."

Leo: Were they glad they bought these.

Steve: Yeah. She said: "I just wanted to give you a heads-up that we so appreciate your work and [she has in quotes] 'advertising' for us that we are sending you a thank you package." And actually I received it. She sent me two headbands as a courtesy and some gift cards for a meal at a local restaurant.

Leo: Wow, wow.

Steve: So I thought that was very nice.

Leo: Boy, Steve. You're lucky you're not in broadcasting. I'd have to bust you for plugola now. Yeah, yeah. You got a little side bet here. All the apple garden you can eat.

Steve: She says: "We know we do not have to, but we wanted to, and show our appreciate for your like of the product and letting people know about it. We did not raise the price as we wanted you to have credibility with your listeners and give them a great deal."

Leo: Nice.

Steve: "Thank you very much for the broadcast, and we are trying to make all your listeners happy. We do not usually ship out of the U.S., but you have a few followers in the U.K., Canada, Brazil, and Australia that would like them."

Leo: Holy cow.

Steve: "We are making special arrangements for these followers."

Leo: Wow. Wow.

Steve: Then she says: "Thank you again. I would like you to know we are a team here, and with Angie, Jose, Chad, and Angela's help, we are getting these packages out in a timely manner. Dawn Ames."

Leo: Wow, nice.

Steve: Anyway, I'm impressed with them. I kept worrying that they were going to jack the price up because they've got something that - they sold a thousand of them last week. They made $40,000 last week. And they may imagine they could make it $50 each, but they haven't. So again, if anyone didn't look, the GRC.com page last week was just a link to their auction. I now have had time and had a lot of feedback from users, a whole bunch of first-use tips. So if you have any trouble at all with it, GRC.com/zeo is now a rather comprehensive page to get you started using your Zeo. And people are beginning to play with it and actually are loving it a lot.

I did want to mention that "The Expanse" is over. I watched the final two episodes. I didn't see them, I think they're on Tuesday nights, so I didn't see them by this podcast last week. But it's just fabulous. And I found myself actually wishing for an inversion. Normally I read the books because the books are so much better. In this case I found myself wishing that I had not read the books so that I didn't know what was going to happen because the presentation was so good.

Leo: Well, good news. I haven't read the books.

Steve: Good.

Leo: I think I started "Leviathan Wakes," but I don't know if I finished it. It's gone beyond that; right?

Steve: Well, we only get 10 episodes. And I think the last episode - oh, and I forgot to mention. Remember how I've been saying how unhappy I was with the beginning? That the very beginning seemed really wrong? That no one who watched it would have any idea what was going on?

Leo: I was puzzled, yeah.

Steve: It turns out they did this out of sequence. They took us back to the beginning, having teased us with, like, what happened after that, and completely filled in all of that missing stuff so that it's now completely clear.

Leo: Should I watch it in a different order?

Steve: No, no, no. Do it their way. I'm always annoyed when shows do that. That's something that people do now that I find really annoying is you've got to keep track of, wait a minute, did this happen now? Who knew what, when? But still, it does bring us current. And I'm getting a lot of feedback from other people who are saying, yeah, this thing is great sci-fi.

Leo: Nice. Can't wait.

Steve: I did want to mention two recent breakthroughs in energy. On the 10th of December, so early last month, Germany fired up what they're calling their "Wendelstein 7-X," which is a funky type of tokamak which is actually called a "Stellarator." And it has seven weird twists in it. So it's as if - I'm sure there's some sort of - it reminds me of some sort of a doughnut of some sort, sort of like an old-fashioned, except if you took an old-fashioned doughnut and kind of did a Mobius thing to it so that it has seven twists. It turns out that a supercomputer designed it. No human could have done this without a supercomputer. But the nature of the bizarre asymmetries they believe gives them a tremendous advantage. So it was able to attain 80 billion degrees for a quarter second. And of course, just to remind people...

Leo: Eighty billion degrees?

Steve: Eighty billion degrees.

Leo: That's, like, off the scale. I mean...

Steve: It's very hot.

Leo: It's hotter than the sun.

Steve: Well, and we need that for fusion because we don't have gravity. The sun has the advantage of a huge amount of gravity.

Leo: Oh, a lot of pressure, so...

Steve: Yes. And so we're trying to do this in a basement somewhere, instead of the center of the sun. But what caught my attention was a quarter second. That's a long time. Because these things are normally, like, okay, we're sure that happened, but it was so short we need to make sure that the recorders were running. But 80 billion degrees for a quarter second. Everyone was jumping up and down. Now, unfortunately, well, not unfortunate, I don't mean to be - I don't mean to say that. China yesterday did 50 million degrees for 102 seconds. Oh, is that a picture of the Stellarator?

Leo: Yeah.

Steve: Yes. Isn't that the funkiest looking thing you have ever seen in your life, Leo?

Leo: Only in Germany would you think that this is a good idea.

Steve: Oh, I know.

Leo: What the, what the, what?

Steve: If anyone's curious about fusion, there are some wonderful animations there about, oh, my god, you just look at it, and it just makes your head hurt. It's like, what in the world? Look at the shape that the...

Leo: How do they heat things up? What is - geez.

Steve: Yeah. It's just crazy. It's funny, too, because the support system, the support infrastructure for the thing is like seven times larger than it is. It's this little crazy-looking doughnut thing, all funky and twisted and bent. But then everything it needs in order to feed it is just like, way bigger than it itself. But anyway, China has built a traditional...

Leo: This is the Max Planck Institute. So this is credible. This is not some...

Steve: Oh, yeah. Yeah, yeah.

Leo: Yeah, yeah. So, but China is trying to do it, too, huh?

Steve: Well, they went to 50 million degrees, but for a stunning 102 seconds. Which is, like, amazing. Now, some of the reporting, I said, let's remember this hasn't been independently vetted, blah blah blah. But, okay. Yes, you're showing the animation of how the seven twists in the Stellarator works. Anyway, that's a great video.

Leo: It seems made up.

Steve: Doesn't it? It's just gotten too crazy.

Leo: If this were in a science fiction book, you'd go, "Oh, come on." Apparently they built it in Minecraft.

Steve: Yeah. The graphics are a little...

Leo: Primitive. Primitive, let's say. But, you know, they're busy doing other stuff.

Steve: Yes. Actual...

Leo: Oh, my goodness.

Steve: Look at that sucker. Yeah. And, like, nothing is symmetrical, nothing - it looks like something...

Leo: Alien technology, it looks like, is what it looks like. It's bizarre.

Steve: Yeah, it totally does. Yeah, and so that shows you how, like, all, like, individual, bizarrely designed, computer-designed rings. And in order to have instrumentation they had to have holes in different places where it would be useful to see something. And so there is an example of one of the magnets, and another one, and another one. Every single one...

Leo: So they're using magnets. Is it particles they're accelerating? Or what are they...

Steve: It is a hydrogen plasma.

Leo: Oh, they're exciting the gas.

Steve: Yes. So it's a plasma. And so you need to heat it up and squeeze it. And so it uses - the plasma is conductive. And being conductive means that you can influence it from an externally applied magnetic field in order to contain it. And so all those crazy magnets are, like, supercomputer-designed in order to move the plasma through this conduit in the proper fashion.

Leo: So of course the problem with this is you put so much energy into generating...

Steve: Right, right.

Leo: ...the conditions for fusion, that you're not getting a net gain.

Steve: Right. And of course - and there are some now that are beginning to be break-even. And here we're now doing a travel, a trip through the inner workings to show - and those are equal potential lines of magnetic force that we just flew through.

Leo: Well, of course they are.

Steve: So cool. And so...

Leo: So are they getting to the point where they have enough, they're getting enough energy out to justify the energy in?

Steve: Yes. We're not quite at breakeven, but we're beginning to be at the point where we're not having to beg for money to pay the electric bills anymore.

Leo: Oh, well, that's interesting.

Steve: Where it is beginning to actually work. And these are all early proof of concepts. But, you know, the concept here is we're burning matter, like the sun does. We're going to simply, instead of fission, which is such a messy energy-generating technique, we're going to fuse matter. We're going to burn matter.

Leo: Instead of breaking down large molecules, like uranium 235...

Steve: Yes, in the same way that the sun...

Leo: We're going to build up hydrogen.

Steve: Exactly.

Leo: From hydrogen to helium or something.

Steve: Exactly. Exactly.

Leo: And energy is released. Step 3, profit.

Steve: Yeah. It's funny because the person who tweeted the note about China to me this morning, he sent back, "Yeah, but we're a long way away from having that in our kitchen or pocket." And I said, "Yes, but remember the size of the first computers."

Leo: Right, right. It may not be as long as it looks.

Steve: Yeah.

Leo: The first thing is to make it happen.

Steve: Yeah.

Leo: Have they actually had a fusion occur?

Steve: Yeah. Sustained fusion reactions.

Leo: Wow.

Steve: Exactly.

Leo: Wow. So they're taking deuterium and turning it into tritium or something.

Steve: Yup.

Leo: Adding a proton there. Wow.

Steve: Boom. It's hard to know if one of those things is out of place. That's just crazy.

Leo: I don't know. It's just amazing. This is amazing. It's crazy.

Steve: So I had a nice tech support interchange that I wanted to share, and also solve a problem that some SpinRite users have. Greg forwarded me part of the email thread, but not the person's first contact. He started with his reply, where Greg says, "Do you have an internal SATA drive?" So the guy probably said SpinRite won't run. I'm getting this red screen that says divide, like, I can't remember if it's divide by zero error or something. That happens to some systems on some motherboards.

So Greg replies, "Do you have a internal SATA drive? If so, in your motherboard BIOS settings you'll likely find an option to switch your SATA motherboard to something called Legacy IDE, Legacy Operation, or Compatibility Mode, maybe just ATA or IDE." And Greg says, "Various BIOSes call it different things, but the idea is that the SATA controllers are made to appear as standard traditional IDE ATA drives. We've encountered instances where the newer SATA BIOS code was less well tested than the older and more solid Legacy code which SpinRite uses. You should restore your motherboard to its previous SATA setting once you're done running SpinRite on your SATA drives."

And so that's the case. What's happened is that, by the time that the AHCI technology came along, no OSes were really using it. They would just use it to boot themselves, and then they had their own AHCI drivers. And that's the support that I most recently added to SpinRite as we moved to SpinRite 6.1. So I'll be solving this problem by moving us away from using the BIOS. But at this point SpinRite 6 still does. The problem is that

SpinRite uses the heck out of the BIOS. And if you have a late model drive, the BIOS defaults to AHCI mode, or SATA mode, using BIOS code that is buggy.

So, and this has been the conundrum is it isn't a bug in SpinRite. There's nothing I can do to fix people's motherboards, except to cause SpinRite not to use the motherboard at all. That's the main feature of 6.1 from which we'll get all kinds of benefits. In the interim, simply switching it back to ATA uses the Legacy BIOS code that does not have the bug. And people think, oh, isn't that slower? No. It turns out it's not any slower than using AHCI. You're still able to drive the drive at full speed.

Anyway, so in receipt of this email, someone named Forrest said, "Greg, thank you. I'll make the suggested changes to the BIOS and start SpinRite again."

So then the following email from Forrest: "I wanted to let you know that the BIOS ATA settings allowed SpinRite to run, and it worked. I was able to recover the pictures my daughter had taken around Europe in the fall of 2014 when she was doing a study abroad her junior year in college."

Leo: Aw.

Steve: "Now she'll be getting the pictures for her Valentine's Day present. She studied marketing overseas; and, as a photography major, she was very disappointed when the last week of her semester the hard drive crashed. I had been saving up money to send the drive off to a data recovery group. But thanks to SpinRite, that wasn't necessary. Now she can add a few of her pictures to her website. Thank you for creating such a great tool."

Leo: Wow, Steve. That must make you feel pretty good.

Steve: Yeah. That's...

Leo: Wow, what a nice story.

Steve: That's the payoff.

Leo: Isn't that great. All right, Steve.

Steve: Forgot to mention one thing. It's in my notes, and I just skipped over it. The Zeo Sleep Manager pods only run with Android.

Leo: Oh, you have to have an Android phone. Because they for some reason left the software on the Android store.

Steve: Correct. And Amazon this week, through the 13th, through February 13th, took $10 off of the seven-inch Kindle Fire tablet, so it's 39.95, $40.

Leo: And you can use that with it?

Steve: Yes. I do, and several other people are using it successfully.

Leo: So it's not just in the Google store, it's in the Fire store, the Android, or the Amazon store. Because that's a different store.

Steve: Actually, no, it's not. You need to run a patch. I have the links to it on my page. But someone's worked out how to install the Google Play...

Leo: Do a side load, okay.

Steve: Well, you can side load, or you're able to run Google Play, the Google Play services on your Kindle Fire and then have access to all the Google Play assets, including the Zeo Sleep Manager Pro app.

Leo: That makes it much more attractive.

Steve: So, yeah. And that's what people are doing is they're iOS users, so they don't have an Android, but they just buy a cheap Android tablet. And I just wanted to mention that this week Amazon has them for $40.

Leo: Nice.

Steve: And it works great.

Leo: Cool.

Steve: So, okay. Nate Smith tweeted me. He said: "Hi, Steve. Could I open up some ports on the IoT router and reach in from within the private LAN? Seems like NAT on the root router would push a request over to the IoT modem's WAN IP. Does this break anything?"

And this was sort of my opening, and a great opportunity for essentially giving our listeners one additional toolkit for ways to understand this so that this and other sorts of problems can be solved. As I mentioned last week, the way IP routing works - so basically we have IP routers. Within each network is Ethernet as the carrier medium for these IP packets. Yet, as we said, the Ethernet is locally constrained.

So a router has two sides, the WAN side and the LAN side. There's an Ethernet network on each side, but yet it is an IP router. So the Ethernet serves only to get the IP packets to and from the router. The router only deals with IP. So we get Ethernet isolation. And that's the key of the security to our three-router solution, the "Y" approach, is that there is no bridging or communication of Ethernet across the routers. There is only IP routing.

And so what I want to focus on a little bit is some subtleties of IP routing in this context. So we all know, we've seen 192.168 dot something dot something. Or maybe 10 dot something something something and so forth. And there's this concept of blocks of IPs being in the same network. So, for example, 192.168.0.1 through .255, that is, where the last tuple, the last byte of the IP address numbers individual machines in the network.

And the key here is what does "in the network" mean? Well, "in the network" means that they all share the 192.168.0, that is, they all share the prefix and then have different suffixes. And the way the prefix is determined is the so-called "subnet mask." That's that thing that's 255.255.255.0. And that number, the 255 is all ones in binary. So eight ones for the first byte, eight ones for the second byte, and eight ones for the third byte, which correspond to 192.168.0.

So what the subnet mask tells every device in that network is any IP where the bits are all the same as its own, wherever the subnet mask is ones, is the same network, that is, the same local network. And that's crucial because routers, when they receive packets, we've often talked about how router receive packets, look at the destination IP, and then decide which interface to send the packet out of. Well, that's also happening in our homes, with our own home routers, exactly like big iron routers out on the Internet.

So the idea is that, with this three-router "Y" configuration, there are three routable IP networks. There's the Y-configured IP network which interconnects the three routers. And that's the LAN of the root router that connects to the Internet, and the WAN ports of the two isolation routers that forms the "Y." And then the two isolated LANs are on the LAN side of at least two, maybe three or more routers, however many you want, as we were talking about.

So here's the tricky part. Those intermediate routers, the ones that connect to the private LANs, they will only function if there's a different LAN on each side. That is, in order for packets to cross, they must arrive at the router and be destined for somewhere outside of the router so that the router sends them out of its WAN link. So, for example, this means that, when you connect the interior routers to the exterior, to that root router, their WAN IPs are going to be assigned from DHCP, whatever that root router uses as its default LAN. Let's say for argument 192.168.0 is that root router's network. I don't know, and it's really a function of the router's firmware, what a standard consumer router does if it receives that network on its WAN. That is, normally consumer routers default to 192.168.0.* or .1.*.

Now, if the router had dot zero dot something, dot one, on its WAN, maybe it's smart enough to choose a different internal network for its LAN number, but it may not be. So that may be something - this is always user configurable. It may be something you must configure, that is, one of those internal LANs you could set to. And for some of this more advanced inter-private LAN routing you're going to want to, you might deliberately set that to 192.168.1.0 and set the other internal private LAN to .2.0. You can generally change them to be whatever you want to. So that clearly and uniquely numbers each of the LANs that you've created. Dot zero dot whatever is the interconnection LAN for the routers. Then one router, whether it's your private or your secure internal private LAN, that might be dot one dot whatever; and the Internet of Things could be dot two dot whatever. And deliberately setting the numbering both assures they're not in conflict and lets you know what they are.

So Nate asks, that sort of started this whole thing, "Could I open some ports on the IoT router and reach in from within the private LAN?" And the answer is yes. The idea would be that you would know what the WAN IP of the IoT router is. And in fact it's easy to use DHCP, the Dynamic Host Configuration Protocol, because it just establishes IPs for you. For this sort of a network, you may want to just assign these IPs manually, even on the outside of those routers. We were talking about assigning .1.0 and .2.0 on the inside LANs. You may want to give them - you may want to assign their WAN IPs also manually. So it'd be like .0.1 and .0.2 for the same routers, respectively, on the LAN side. So that you know how they're numbered and so that their numbering is not changing.

So the idea is that any ports that you open on those interior isolation routers are addressable by that router's WAN IP. So, for example, if you are in one of the private networks, and you know the IP and port number of something that you have deliberately port-mapped through the other router, you can address it using the WAN IP of the other router and the port number that you have mapped through. So it's necessary, to do this kind of stuff, you sort of have to take things off of autopilot so that you know how things are numbered. But the concept is each of the networks should have its own IP space that are non-overlapping.

The reason I mention that is, technically, nothing would prevent both of your interior routers, your IoT and your secure router, from both being 192.168.0. That is, they technically - or dot one. They have to be a different network than the "Y," the router interconnection network. But they don't have to be different from each other, strange as that sounds, because even though, well, it's very much like what we have right now on the Internet. Everybody in their own homes has 192.168.0 dot something. And your next door neighbor has 192.168.0 dot something. That is, we all are using overlapping networks, yet there's no confusion because the NAT'ing converts the IP to the public IP. That's what's unique. The networks behind the routers don't have to be.

So just that same principle applies in our home network if we've got a "Y" routing or a multiple router with network isolation system. They don't have to be numbered separately. But if you don't number them deliberately and separately, you cannot route between them because, if they were both 192.168.0 dot something, then any packets you were trying to send from your network - wait a minute, no, that's not the case. If you were, yeah, because you were addressing to the other router's WAN IP. So you can number them the same. Sorry about that. I confused myself. This is a little confusing. As I'm sure everyone now knows. You are able to route packets because you are addressing to the other router's WAN IP, which you probably want to make sure has not been addressed dynamically.

So Nate's answer, the answer to Nate's question is yes. Again, take it off of autopilot, manually assign these IPs if you want to do more fancy work like this. Or, if you just plug things together in autopilot mode, and it sort of doesn't work, it may be because the routers just don't know how to handle a 192.168-style IP on the WAN side. So you may have to manually configure either the WAN side to be different from the LAN side, or the LAN side to be different from what it received on the WAN side.

So the second question I pulled, I sort of may have covered. Ken Hundley wrote, he said, "Hi, Steve. Your 'Three Dumb Routers' podcast answered several questions I had about this, but I have one more that I'm hoping you can answer and may be something worth mentioning on your show. I have a fourth router that I need. It's required by Comcast for their Home Security system, and I have no access or control over it. With this setup, I would have at least three wireless routers. Will I run into any potential issues having so many wireless networks in such a small area?"

And I know of none. There's really no practical limit to how many routers you can have. I mean, "practical" meaning up to 254 or so. But certainly three or four shouldn't be a problem. And maybe, if all of the routers were next to each other and saturating their bandwidth and all on channel 11, that could cause some problem. But probably your IoT devices have very low bandwidth. I don't know what Comcast and their home security system is doing, or maybe whether it's also being used to stream media.

But it makes sense to physically move the access points from each other a few feet apart, if you can, and maybe to put them on different channels, just to give them their own WiFi channel space so they're not really actively fighting each other. But otherwise they ought to work just fine. And you can run as many routers off of spokes from the root router as you want to.

And finally, Yuki Takizawa said, "Hey, Steve. I just listened to the 'Three Dumb Routers' episode. So I just wanted to DM you a quick question. I get the isolation property with any of these approaches. My number one issue is sharing a printer on my home network among secure and guest networks. Perhaps a more sophisticated VLAN routing software/configuration could do this. But it seems hard with a simple setup, even if I use multiple routers."

And to Yuki and everyone else, it's really not necessary. This is covered by what Nate was asking, that is, you could leave the printer behind one of the interior routers and map a port through which the other network would know how to get to. It is the case that packets from one network would then be transiting from the shared "Y" network into the network where the printer is. So, for example, if the printer had known vulnerabilities, and it was in the secured network, something malicious could potentially find the printer and crawl into the printer and get up to some mischief. But you could put the printer over in the untrusted network, on the IoT network, and then map a port to it so that from your trusted network you are able to reach through both routers and into the printer. Or you could put the printer out in the common network, that is, it would be on the IP network that all of the routers share, that links them together, and then be available easily to any of those.

Again, though, the danger is that - and this is why we have this network isolation behind routers. You wouldn't, if that printer were compromised, if some malware somehow got into the printer, it's then on the common network, potentially, thanks to it being Ethernet, able to see all the traffic coming and going. So I think I would put it over in the IoT side. I would treat it like an IoT device that I want to give some selective access to from my secure network. And so it just looks like the IP of the IoT WAN. And any traffic that goes out of my secure network will be routed over through the other router and to the printer.

So you absolutely can selectively create controlled intercommunication between the secured networks. And as Nate indicated, and Yuki gave us an example, there are reasons why you'd want to. So that's how it all works. I think we got this one taken care of.

Leo: Yeah. All of your questions answered to your satisfaction, one hopes. But of course there's always next week.

Steve: And there's tinkering and playing with configurations, and a lot of people liked this idea. It was way more popular than I thought because it does offer such good security against this growing threat from wacky devices where security is an afterthought. And I'll probably share what Brian Krebs wrote about this, and some of his thoughts, next week when we talk about it. I'm sure this is where an increasing percentage of the podcast is going to go. So this is one of the reasons I wanted to take some time. Even though we have talked about it before in less detail, I think creating a truly isolated network for this crazy interconnected device stuff makes sense.

Leo: Steve, once again, you've done it all, in a mere two and a half hours. If you want to watch the show live, bring lunch and tune in every Tuesday at about 1:30 p.m. Pacific, 4:30 p.m. Eastern, 21:30 UTC - you could bring dinner, too - and listen in. We do the show live then. But of course always, after the fact, on-demand audio and video is available. Steve has it at his site. That's GRC.com. While you're there, by the way, lots of free stuff, including SpinRite, the world's finest hard drive maintenance and recovery utility; his sleep research; his Vitamin D research; SQRL, which is going to change the world of logins forever; and much, much, much more. GRC.com. He also has transcripts, written transcripts there.

We have audio and video at our site, as well. That's TWiT.tv/sn. There's a YouTube channel. You can find that at YouTube.com/twit. Look for the Security Now! channel. And of course you can subscribe on your favorite podcatcher because after 11 years there's not a podcatcher around that doesn't have this show. Security Now! on the TWiT network. Thanks, Steve. We'll see you next week.

Steve: See you next week, my friend.