Listener Feedback #218

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

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-524.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-524-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We are going to get some questions and answers in, finally. We'll talk a little bit about adblocking, yes, but also about the Windows Patch Tuesday. Today's the day. Security updates from a lot of vendors. And a kind of hard to believe flaw in Seagate's hard drives. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 524, recorded Tuesday, September 8th, 2015: Your questions, Steve's answers, #218.

It's time for Security Now!, the show where we cover all of the security news. This is the super geeky show, frankly, on the network because we cover anything that Steve's into. And since he's a super geek, that could be anything from Vitamin D to BSD routers and everything in between, including great science fiction. Hi, Steve Gibson.

Steve Gibson: Hello, my friend. Great to be with you again, as you get ready to wing your way to the East Coast.

Leo: Yeah.

Steve: For your meet-up. And I'll be watching this channel tomorrow as Apple unveils their next set of updates.

Leo: I don't know what I'm going to do. I guess I'll be in meetings. I can't imagine.

This is an example of me being in the real world for a change. What do people in the real world do during an Apple event? Do they - seriously. I'll be at SquareSpace. We've got - we're meeting at SquareSpace, and I've got to figure they're interested in this stuff.

Steve: Well, and because for the last 10 years you've had TWiT.tv and the TWiT Network, and you've always been surrounded by your expert panel, who are watching this stuff all happen in real-time.

Leo: First Apple event I've missed since the iPod in 2001.

Steve: Well, we all know that it doesn't really happen until ordering time, midnight on Friday we're presuming.

Leo: Or 3:00 a.m., if you're on the East Coast.

Steve: That's right.

Leo: Which I am. Ay ay ay. Anyway, I will watch with interest. Are you going to get the new - a new iPhone, do you think?

Steve: I really do want it. I do. I want, I mean, I use the iPhone as my go-to device. My iPad I use more than any other computer in my life is my iPad, just my lifestyle. You know, I take it with me when I leave the house and relax when I'm having a meal and read stuff.

Leo: Would 12.9 inches be too big for you?

Steve: Yeah, you know, I even got the mini, the latest version of the mini, and I actually returned it because I thought, eh, you know, there was something I was able to say it wasn't doing right. I don't remember now what it was.

Leo: Ten inches is just the right size for you.

Steve: I really think, yes, the standard iPad is just fine for me. I don't know how they fix that. I don't think I need a big one. But I would love to have a stylus. I don't know why.

Leo: Yeah.

Steve: But I just sort of think that's a cool thing. I've, you know, the idea of being able to jot a note or twiddle or doodle. And maybe Apple will do a good job on that, rather

than giving us something that doesn't really work like they did with the Apple Watch.

Leo: One thing you can be sure is that they will not have a stylus that goes in the hole the wrong way.

Steve: No.

Leo: I don't think Apple will do that to us.

Steve: So we have a Q&A today since the world has been kind to us with news. We have some interesting things to talk about, but not - we're not overwhelmed so that that alone will take up the whole podcast. So we've got a bunch of great questions. Naturally, lots of stuff from our listeners following on from the discussions we've been having.

So Seagate suffered a surprising problem with a WiFi hard drive which I just - when you read - when I explain this to people, they're just going to put their head in their hands. It's like, in 2015, how can this still be happening? Adblock has released an Adblock Browser, just this morning, of course on the eve of iOS9. We'll talk about that a little bit. There was some weird belief that suddenly Chrome was trying to defeat adblocking on YouTube, which turned out to be specious. Android phones have been...

Leo: uBlock Origin works on YouTube.

Steve: Yes.

Leo: Yeah.

Steve: Android phones are being shipped with preinstalled malware. And I wanted a little bit of an update on...

Leo: Oh, that's convenient.

Steve: Yeah. Users don't have to install their own.

Leo: Nice.

Steve: Yeah. And some feedback from my click-to-play recommendation. So, and then, of course, 10 great questions from our listeners. So I think we have a great podcast in store.

Leo: Jam-packed episode. And good news, my flight doesn't leave till 9:30. So

plenty of time.

Steve: Did you make it to the DMV last week?

Leo: It was - thank you. I did. It was great. In fact, it was amazing. They let you make appointments. I had an appointment for 3:15, got there at 3:10, waited in line for, like, eight minutes. I thought, oh, I'm dead, because the line was so long. But they said, no, no, you're on here. And within three minutes I was out the door.

Steve: Nice. Wow.

Leo: But the line was three times longer than the actual stuff. It was great. Make an appointment. Steve Gibson, Leo Laporte, Security Now!. Let's get into the news.

Steve: So our picture for the week ties into one of the questions that a listener asked about his puzzlement over glass platters because he was...

Leo: Yeah. I was puzzled, too, until I almost my eye poked out by one.

Steve: Actually, they're really dangerous.

Leo: Yeah.

Steve: If you poke around the 'Net much, there's like people saying, oh, my god, don't, you know. I guess one of the ways people are destroying drives is they're using some sort of a device to just push right through the axle of the drive, like some sort of a punch. And if you do that with more recent glass or glass-ceramic platters, which we'll be talking about later in the show when we get to this guy's question, I mean, they shatter into microscopic shards. And I saw one person saying don't ever do that over carpet, or you will never get the carpet cleaned of all of the glass that is in there. So anyway, I just saw this fun picture that showed, yeah, this is not a platter that bends.

[Video plays]

Leo: All right. Get ready, because I'm going to show you the video of Patrick.

Steve: Oh.

Leo: That's the drive. He shatters it into a million pieces. Watch, let me show you again because he just misses my eye. So he didn't know. He thought it was metal. That was a few years ago, Steve.

Steve: Wow.

Leo: So this has been around for a while.

Steve: Oh, it has. And when we come to the Q&A we'll talk about why that's the case, and actually why it's...

Leo: Why glass?

Steve: ...where we're headed in the future.

Leo: Oh, yeah.

Steve: Yeah.

Leo: And wear protective wear if you decide to destroy it with a hammer.

Steve: Wow.

Leo: Yikes.

Steve: So this is the second Tuesday of September. So we are at...

Leo: Patch Tuesday.

Steve: Patch, thank you, Patch Tuesday. And no earth-shaking news. There were 12 update bundles, five of which Microsoft rated as critical remote code execution exploits. IE and Edge both get updates. Now, Edge, of course, now that we have another browser from Microsoft in addition to IE, it's getting critically updated also. The one worrisome thing is this so-called "graphics component," they said, which affected Windows, Office, and Link. So that's sort of scary because that tends to be in the kernel. And if history is any teacher, just rendering a specially malformed image can be all that is needed in order to get a remote takeover.

And then Windows Journal, Office, Media Player. Hyper-V had an update that said "security feature bypass," which is, you know, never what you want to hear in a VM manager. So I would say to all Windows users that we are Second Tuesday of the Month. Update Windows as soon as you can. I thought I saw - mine came in after I was already up and running and had Skype up. I though, oh, I'll wait till after the podcast because I run Skype on a Windows 7 machine, and so it's getting updates [crosstalk].

Leo: And I'm doing something that you would never do in a million years. I am not only running Windows 10, I'm running a beta of Windows 10. So, yes, let's see. Windows Update, yes, indeed.

Steve: You're on the fast loop of Windows 10?

Leo: Fast, fast, yeah.

Steve: Yup.

Leo: Yeah, yeah.

Steve: So, okay. Seagate. Believe it or not. This drama began quietly, because this was responsible disclosure, back in the middle of March of this year, March 18th. A company named Tangible Security that we've never run across before, they found just a shocking problem with Seagate-generated WiFi drives. Seagate has these wireless hard drives which they sell under the Wireless Plus Mobile Storage and then just Mobile Storage names. And then La Cie - is that how you pronounce it? L-A C-I-E? They are a relabeler. Their version of the same drive is called the FUEL. So, and it may well be that others that have been OEM rebranded of the Seagate Wireless, you know, these WiFi drives could have the same problem.

Well, it turns out, it's just hard for me even to believe this, that in the firmware of this wireless hard drive they undocumented and hardcoded remote Telnet access with the default credentials of root as the username, and then the device's default password. And it's like, no, no, no, no, no. What year is this? This is 2015. So what that means is that anyone who has WiFi access to the network that this thing is on can Telnet to this wireless hard drive. Which is to say, you know, our techie listeners know that basically Telnet is a remote command prompt. It's a remote console.

Leo: Shell, yeah.

Steve: Yeah. So...

Leo: And nobody uses it because it's insecure in itself.

Steve: Oh, it's like, yeah.

Leo: Sending it in the clear.

Steve: Everyone now uses SSH, which is an encrypted - it's sort of like SSH is the SSL version of Telnet. But so this is an unencrypted, in the clear, what is it, port 23, I think?

Leo: Twenty-two? Twenty-three? Yeah.

Steve: Just sitting there, ready to, you know, accepting TCP connections. And so you give your Telnet client, and it's like, log me in as root and whatever the default password is for the device, that is, the factory password, and you get a prompt. It says, hi there.

Leo: Oh, my god.

Steve: Well, hi there, root. What would you like to do? So...

Leo: That was obviously left in for remote patches and support and administration. But why they would leave it in the firmware...

Steve: Well, yeah. So first of all, you would never want your root user to be called "root." You know, at least name it gibberish. But the problem is we know this is not safe. People are going to look at the firmware, or do a port scan. You can do a port scan of your hard drive, and it's going to be answering TCP connections on the Telnet port, and then that's going to beg the question, oh, well, Telnet is - so that means that the Telnet service is running on this hard drive. Okay. Even that phrase is bad.

Leo: That means you have a daemon running on the software. Like, that's crazy.

Steve: Yeah. Yeah, it's nuts. On a hard drive. It's like, okay. We don't want that. So anyway, the good news is there is an update. You can download it from - so what happened was Seagate was notified on March 18th. In a very short time, I was impressed with this, 12 days later, on the 30th, they replied and confirmed there were vulnerabilities. Then, unfortunately, it took a hundred days before anything happened. So this has been out there since the release of these drives. And these guys found firmware dated October 2014.

So they said: "The following devices with firmware versions 2.2.0.005 and 2.3.0.014," they said, "dating back to October 14, are vulnerable to three" - and I've only talked about one - "three attack vectors." And then they said: "Other firmware versions may be affected, as well." So the takeaway here is, if you have Seagate WiFi hard drives, you want to go and update your firmware.

So the first of the three was their use of hard-coded credentials to give a Telnet user root access to the drive. And these guys wrote: "The affected device firmware contains undocumented Telnet services accessible by using the default credentials of 'root' as the username and the default password. An attacker can covertly take control of the device, not only compromising the confidentiality of files stored on it, but use it as a platform to conduct malicious operations beyond the device." Because of course they can download and run any other services that they want to. I mean, this is just unbelievable. Okay. Second problem, that they called direct request, "forced browsing."

Leo: Wait a minute. There's another one?

Steve: Oh, there's two more. Yeah.

Leo: What a mess.

Steve: I know. "The affected device firmware provides unrestricted file download capability. Attackers can gain access to all files stored in affected devices." This is through some other undisclosed mechanism other than this Telnet problem. So "The affected device firmware provides unrestricted file download capability," meaning that there's no security, basically. "Attackers can gain access to all files stored in affected devices. This vulnerability requires attackers to be within range of the device's wireless network."

Well, yeah, because it's a WiFi device. So maybe that means - so that wasn't a limitation on the Telnet access. So this may be different. This, for example, maybe it's not routable through the border router. Again, we're scant on details because they're not wanting to talk about this until everybody gets this fixed. And right now, this just happened. So right now nobody has this fixed.

And third, unrestricted upload of file of dangerous types. "The affected device firmware provides a file upload capability to the device's /media/sda2 file system, which is reserved for filesharing. This vulnerability requires attackers also to be within range of the device's wireless network in order to upload files to it. If such files were maliciously crafted, they could compromise other endpoints when the files are opened." So, wow. About as bad as it gets.

Just, again, like here's Seagate, a company with a great reputation. On the other hand, we do know that many of these high-profile companies are getting their firmware from third parties in the same way that TP LINK made the hardware for Google's OnHub.

Leo: It also could be a reference platform that they left that in for remote updates...

Steve: We've seen that.

Leo: ...of the software, and you're supposed to take it out before you ship, and...

Steve: Yes. I mean, how hard is it to stop a service from running on boot in Linux?

Leo: Right.

Steve: I mean, it's like not. You just don't. So it's like they shipped it by mistake and left the Telnet service running.

Leo: And of course there is one way it could be worse, if Hillary Clinton used it for her email server. Then we know. Then we know we got "trouble right here in River City."

Steve: In River City, yes. Okay. So in an odd piece of news on today, Tuesday, the day before the Apple iOS9 announcement, where one of the major announcements is, I mean, that's making just as much news as the gossip and rumors for what next hardware they're going to be producing, is the addition in iOS9 to Safari, the default browser, of course, of hooks which will allow adblocking extensions to be created. And we know that there's already one called Crystal, which is highly anticipated for Safari, and uBlock.

I have seen nothing from Raymond and his branch, uBlock Origin. But Chris, the guy who's maintaining the unbranched or unforked uBlock, has announced that he will have a Safari version of uBlock, which is fundamentally the same as uBlock Origin, available soon. So apparently the Adblock Plus people, or the Adblock people, they're a company called Eyeo, E-Y-E-O, which sounds like a nursery rhyme.

Leo: And we're so close.

Steve: Yes.

Leo: If it had just been EIEIO, we could have really loved it.

Steve: Exactly. And everybody could have remembered it. And in fact you need that because, if you look for using Apple's horrific search in the App Store, I just can't...

Leo: Oh, yeah. Can't find anything, yeah.

Steve: Every time I go looking for something I think, how can you not have figured out search? Something like that everybody else, especially your major competitor in the world, has nailed search, but you can't do search in an App Store.

Leo: Yeah.

Steve: Anyway, it's not easy to find this. There are many things called Adblock. So you need to look for "adblock browser from EIEIO." No, "from Eyeo."

Leo: Now you're really confusing everybody.

Steve: That's the one you want. And just for the hell...

Leo: Is that the best one, not Crystal? Because people have been talking a lot about Crystal.

Steve: I would wait.

Leo: Okay.

Steve: I'm talking about it because it's there today. We're thinking, what, it may be like two weeks before we can actually add adblocking to iOS. So today, if anyone wanted to experiment with it, the Adblock Browser, I downloaded it this morning so that I could talk about it.

Leo: Eyeo is Adblock Plus.

Steve: Yes, Adblock Plus.

Leo: Plus. That's not the same as Adblock.

Steve: Correct.

Leo: Yeah.

Steve: And so...

Leo: That's these guys.

Steve: Yeah. So, yeah, there it is. That's the one. You want - well, wait. Is it from Eyeo?

Leo: Yes, from E-Y-E-O. Adblock Plus.

Steve: Okay.

Leo: And this one's GPLed and open source. This is not the company that's selling - or is it the company that - I get so confused.

Steve: Okay, now, no. You're looking for a browser. So you go...

Leo: No, no. I know. But this is the Adblock Plus company. Yeah, yeah, yeah.

Steve: Oh, correct, correct, yes.

Leo: Yeah.

Steve: And so in the App Store it's Adblock Browser. And the icon, it sort of looks like the globe, the world globe is wearing a jaunty little stop sign hat because it's sort of off to an angle, because the Adblock logo is the red stop sign. And so they've sort of got - it's got half a stop sign and half of the globe is the logo. So that's the one you want. And I loaded it, and it works just fine. It uses their EasyList, which is the ad server list that they curate. Yup.

Leo: There's the logo. It is jaunty.

Steve: Yeah, it's a little jaunty little...

Leo: I think it's a stop sign, but I...

Steve: Right.

Leo: Does look like a jaunty little hat, yes.

Steve: Yeah, it's meant to be a stop sign. And so we have that today for iPhone and iPad. And there are some features, you're going to want to dig through the menu because they, by default, they block ad servers using their EasyList. And of course that's the one that everybody else clones. I mean, because it's publicly available. So, for example, uBlock sucks that down, uBlock Origin sucks that down, because it's a great, curated, very up-to-date list of ad servers. But under More Blocking Options, which are all turned off by default, they have Disable Tracking is off.

And I note, Leo, that every time you have brought up, and you've been talking about on your other podcasts the pending Adpocalypse, you know, like what's going to happen when Safari adds this to mobile, makes it so easy for users to do this. And invariably one or two of your talking head guests will not, like, I don't really mind ads, I just get creeped out by tracking. And that's how I feel, too. It's the tracking which we don't see. The ads are like the part of the iceberg that's above water. And tracking is, like, way bigger in terms of what annoys people, the idea that, in fact, one of your guests on Sunday on TWiT, two days ago, said that he was a little creeped out when he, like, left one site and went to a different site, yet the ads that he then saw sort of followed him, like from what he'd been doing over to where he was.

And he said, "That creeps me out." He said, "I'd rather have random ads and not have it so obvious that something is following me around the Internet." So under More Blocking Options for this Adblock Browser, Disable Tracking you need to turn on. Disable Malware

Domains is off by default. Turn it on. Why would have that off? So turn that on. Disable Social Media Buttons. I would say yes because that's of course a big - and that's how Facebook tracks people all over the place is all those little Like buttons everywhere. That's sending a ping for your cookie, your Facebook cookie, back to Facebook so they know where you are. So Disable Social Media Buttons. Turn that on.

And then the fourth one I deliberately left off, and that's Disable Anti-Adblocking Messages. And that's because I want to know if a site is unhappy that I have an adblocker because, I mean, I want to support the sites that I visit. I just don't want to get - I don't want to be tracked, and I don't want to get infected with malware in return for seeing ads I'm not going to click on and don't care about. So anyway, dig around in here. There is, you know, the whole Adblock Plus deal is their so-called "nonintrusive ads." And so the other option is Acceptable Ads, and they say "Allow some nonintrusive ads." And that's on by default because that's their - that's Adblock's own monetization model, which has been very controversial.

Leo: Yeah, undermining everybody else's monetization model.

Steve: Yeah. And Google and Microsoft and other major players have paid Adblock for exceptions to their EasyList list. And so what this says is, no, no exceptions. I want full blocking. I don't want to accept your monetization model because that's not mine. So anyway, we do have this Adblock browser available today. And it'll be interesting to see how it goes. Ultimately, I think that what we're going to see is people, I mean, I want uBlock. I want uBlock on Safari. I mean, I'm still using Safari as my default browser on my iOS devices, even though it would be easier to use the LastPass Tab Browser. I just sort of like using the native...

Leo: Yeah.

Steve: You know, the one that comes with it. To me that feels better. And so I'll definitely be - and we'll be covering the extensions that are available for iOS9 as soon as iOS9 becomes available to us.

Leo: This is available for Android, too, the Adblock Browser.

Steve: Yes, good. I'm glad you mentioned that.

Leo: In most cases what they're doing is they're basically using the facilities of the built-in browser at WebKit or whatever, and just kind of...

Steve: Oh, yeah. They're not writing a browser from scratch.

Leo: They're skinning it.

Steve: Right. They're not writing a browser from scratch. And in fact in this case they

didn't. There's something called - it wasn't WebKit. It was something kit that these guys used as the armature for their add-on. And of course the fact that this is now in the Play Store is kind of big news because...

Leo: Apple's turned this off before; right?

Steve: Well, it was Google.

Leo: Oh, Google did.

Steve: What happened was Adblock Plus tried to do this in the middle of March, and Google rejected them and kicked them out of the Play Store. The EFF got all huffy about it and did a posting at the time saying "Google takes the dark path, censors Adblock Plus on Android." And EFF wrote: "In a shocking move" - okay, well, I don't know why anyone would...

Leo: Shocking.

Steve: Shocking. Google...

[Crosstalk]

Leo: [Indiscernible] going on.

Steve: "Google has recently deleted Adblock Plus from the Android Play Store. "This is hugely disappointing," wrote the EFF back in March, "because it demonstrates that Google is willing to censor software and abandon its support for open platforms as soon as there's an ad-related business reason for doing so." And then they said: "Google's stated reason for the ban is that the Android app allegedly 'interferes with or accesses another service or product in an unauthorized manner." Which, you know, is corporate legal speak for nothing. You know, we didn't like...

Leo: But doesn't Google actually pay Adblock Plus to be part of this?

Steve: Yes, they're one of the people...

Leo: So confusing.

Steve: ...that sponsor Adblock Plus. Yeah. So anyway, so it'll be interesting to see how this goes. I did run across an interesting site when I was digging into this. And I think it's called - I just tweeted about it. It's Fair Play? No, Fair Page. It's either Fair Page or...

Leo: PageFair. Yeah, yeah.

Steve: PageFair.

Leo: These are the guys who said, as it turns out maybe somewhat inflatedly, that $21 billion would be lost.

Steve: Yes, 22 I think is their number.

Leo: Yeah, in this year to ad blocking.

Steve: Yeah. And what was so interesting is I spent, in fact, I was using the Adblock Browser to browse their site. I thought, okay, this'll give it to them. Check out this user-agent, suckers. And I went through their FAQ because I was just sort of curious, I mean, I'm really wanting to understand the ecosystem of this a bit better. And so their blog is a series of closed but expandable, you know, click-plus-to-open-this questions that they've asked themselves and then answered. And you go through the entire thing, no mention anywhere about tracking. I mean, it was really very fair. Because of course they sort of have a different model.

Their deal is that websites can sign up with them to find out how much money they're losing somehow. That is, so they put some instrumentation on a website which will be adblocker sensitive. It'll be blocked by adblockers. And so, and then they must have some other instrumentation that isn't so they can look at the delta of what got blocked and what was allowed and tell sites how much revenue is being lost. And then they apparently sell sites a service which provides the message that, oh, this site is ad-supported, which appears if you visit the site with an adblocker running, to make the explanation and the plea for turning off your adblocker. So that's sort of their angle.

But they really do a nice job in the FAQ of explaining all of this, except nowhere, nowhere is the elephant in the room, which is the tracking. That's the part that people, from everything I have heard, and like listening to your other podcasts, that's what upsets people. Not the ads as much as the idea that we know that they're trying to monetize us by profiling. And so if there was a stop tracking, but, you know, okay, then that would be a different thing. But nowhere does that get mentioned on this otherwise very complete page. So I thought, yeah, that's interesting, you know, they're not talking about that at all.

So there was some what turned out to be specious news that - and, I mean, there was a lot of talk about this, some buzz in the last couple days, that people were seeing Chrome and the YouTube app for Chrome not blocking ads, even though customers had adblockers. And so first there was the conspiracy theory that, look, oh, look, here it is, this is what we were waiting for, that Chrome had modified itself to defeat the adblocking extensions to allow its own properties not to have ads blocked.

Well, it turns out, just count to 10, and then we'll get an answer. And the answer was there was a mistake. It was introduced by a security fix which at the time was not public. Issue 510802 is a security mistake that said webRequest API allows intercepting XHR from apps and extensions. So XHR is the XML HTTP request API, which is the - it's the

whole Java dynamic web API that allows pages to make requests back to their parent server to create updated content on the fly. For example, Gmail is a heavy user of that. Sort of the next-gen of web is doing all of this.

So there was a security problem, and the fix broke an aspect of the signaling which goes to web extensions - and this also ties into a question that we'll be getting to later - because there's an API called chrome.webRequest.onBeforeRequest, and extensions can register themselves to receive that signal. And as it sounds, it's on before request, meaning send me a message containing the details of a browser request before the browser acts on it. And that allows the extension to examine it and go, eh, no. And based on whether that extension returns true or false, if it returns false, then that request is aborted at that point, and it doesn't go any further.

So what was happening was in a really sort of flaky way some people were seeing this, they were being affected by it; others were not. So it's been fixed. And it may still be in the wild, that is, it may be fixed in internal builds and not yet pushed out. But for any of our listeners who hear about this or may experience it, it's just a mistake that Chrome made, the Chromium project made, and they're in the process of getting it fixed. It's not an exception that Google has made for Chrome that allows them special adblocking circumvention that other browsers or users don't have.

And then the news that, from a German-based cybersecurity firm, G Data, and you may want to bring up - let's see, there's a PDF that I link to lower on the next page of the show notes, Leo, because I've got a couple pictures in the notes. But what they have found, they've done a survey of Android malware, Android smartphone malware. And they have seen a 75% increase in what they call "preinstalled malware" during the past six months. Now, it takes a little bit of digging, but it turns out that all of this is coming from third parties somewhere between the originating manufacturer and the end user. That is, if you buy, sort of on the gray market, if you buy not from someone like a major retailer, like an Amazon or any of the major cell phone carriers directly, but if you get it on eBay from some remarketer person, that seems to be where this is coming from. So they did single out three different manufactures: Huawei, and I guess is it pronounced Xiaomi?

Leo: Xiaomi.

Steve: Xiaomi. Which I've heard, of course, but I never saw in writing, Xiaomi.

Leo: Xiao is little, I think.

Steve: Xiaomi. So, yes, so I now know how to pronounce Huawei.

Leo: Huawei, yeah.

Steve: Now I know how to pronounce - well, I thought I did.

Leo: You're close enough.

Steve: Huawei.

Leo: You know, no one knows how to pronounce it. We're all making it up.

Steve: And Xiaomi. And unfortunately, Lenovo is also - has made the hit parade in this case. But it turns out there's many more. There's a ton of models of Alps Android phones, the A24, 809T, the H9001, the 2206, the PrimuxZeta, the N3, and the ZP100, the Alps 709, the GQ2002 - I don't think I'd buy an Alps phone. There's an Android P8, and then - and it goes on and on. I have all of these in the show notes. There's also the SESONN phones seem to be frequent targets of these.

They said 25 or 26 different smartphone units were discovered to be carrying malicious software before the consumer acquires the device. And the nature of this is unfortunately a little bit stomach-turning. This is infected firmware which knows how to infect the Facebook app and Google's Google Drive app so that, when users install them, if they're not preinstalled, then this firmware is able to reach up and alter the running apps in order to infect them with spyware, and in some cases adware.

So anyway, again, this has been in the news just - it just hit the news. The good news is this is not coming from the original manufacturers of these. Those phones are all clear and clean. It only appears to be when it goes through a third party, sort of the gray market, that these little goodies are being installed.

Leo: Yes. And I can promise you, nobody in the United States who's listening has ever heard of any of these Android phones. These are cheap phones sold in the Third World. I mean, and if you buy a phone, used phone, be careful. But these are - none of these. Do you recognize any of those names? No.

Steve: No, no.

Leo: No. This is not a Samsung Galaxy S6 we're talking about here. This is the P8.

Steve: Well, and so these are the reason that total smartphone sales or total Android platform is just huge numbers. But it's a large number of them are these wacky, no one's ever heard of them, off-brand phones.

Leo: Yeah, not sold in the U.S.

Steve: Correct.

Leo: And, no, you're safe if you buy from Motorola or Samsung or LG.

Steve: Right, right.

Leo: They're fine.

Steve: Right, right. I got a lot of feedback from people. Remember last week, Leo, you and I were talking about that setting in Chrome, if you went to chrome://settings/content, then we talked about how Google was changing the default to tighten things up in Chrome, but that I recommended going one better and using the "Let me choose when to run plugin content." What Google was doing was they were going to start restricting.

And remember that Apple - wait, it wasn't Apple because Apple doesn't run Flash. Oh, it must have been Firefox at the same time where they were going to be restricting Flash from running - I think it was Apple. No, no, Amazon. Amazon was banning Flash ads, and at the same time - no, I'm getting - it's Amazon was banning Flash ads, but there were two browser vendors, sorry I can't remember who, who were both going to be starting to raise their shields, essentially, against running Flash content.

And so the setting in Chrome that I recommended was one better than Chrome was now using, which was "Let me choose when to run plugin content." A bunch of our listeners running Chrome changed the setting, and I received a lot of feedback after, in days following the podcast, that it was working beautifully for them; that is, they were going to sites, and they were seeing blanked out regions with "Click if you want to run this." And obviously these would otherwise have been run, were it not for them choosing this setting.

So I just sort of wanted to amplify that that's great to do in Chrome. And under Mozilla, under Firefox, it's available, but you have to do it again manually. It's still not the default. Go to Tools, then Add-ons under the Tools menu. And then on the Add-ons page select Plugins. There you'll see a list of all the plugins. And over on the far right is a dropdown list box where you can say "Ask to activate." And there you get a weird kind of little LEGO block that you have to click on if you want to run it. And I'm seeing that now wherever I go.

So I just wanted to strengthen that recommendation a little bit more because, with the feedback that I've gotten, people are loving it, just that that's - they're not running Flash anywhere, by default, unless they want to run Flash. And then they are able to click to do so. And you can make site-based exceptions. So, for example, bit.ly. For some reason, they don't have a big Flash presence, but they do some little weird thing where they sort of, when you say I want to copy the bit.ly link, it does a little fading drifting sort of ghost of the link fades away. And I think they're doing that with Flash. So it's like, okay, I can live without that. But, for example, if you wanted to whitelist bit.ly's use of Flash, just so you weren't being asked, you could do that. And then it's a sticky reminder over in Chrome, so you're able to do per-site whitelisting.

And in the spirit of a Q&A today I got an interesting question that I don't - I'm not sure I ever talked about it before, but this is from Quinton is his name, in Nova Scotia, of course in Canada, who said, "How do you test SpinRite?" He wrote: "With the mentions of SpinRite on the last few episodes of Security Now!, it got me thinking about how you actually go about developing SpinRite itself. More specifically, how do you test SpinRite? How do you know that it is detecting an error on the hard drive, and it's not just an error in your code?

"I'm imagining there are two different ends of the spectrum of how it's possible to know exactly where the problems are on a drive for testing and verifying your code. On one

end of the spectrum you're cracking open a hard drive, putting it under a microscope, counting the sectors and bits and strategically destroying some data so that you know exactly where the drive will fail." He says: "On the other side, you have a virtualized hard drive where you're able to corrupt bits at your leisure. How exactly do you go about this? I'd love to hear your thoughts on the next episode of Security Now!."

And Quinton, you get your wish. So there was actually a time, decades ago, where my tech support by my team, back when I had a team at Gibson Research, a bunch of guys doing tech support with me, and development, where we would take the lid off a drive and run the point of a pin radially out from the center of the drive, or maybe sort of in a lazy spiral sometimes, depending upon what kind of mood we were in that day, in order to create absolute physical defects on the drive. These are, of course, I mean, you know, drives weren't that expensive. We had scrap drive. We used drives a lot in testing SpinRite. So some we would deliberately destroy. And of course just taking the lid off a drive allows dust to fly in because this wasn't a clean room environment. So we'd put the lid back on and spin the drive up, and sure enough, we'd have a bunch of problems. I don't do that anymore these days.

But it turns out that there has always been, and still is, a way, using undocumented, supposedly no longer existing, no one knows about them, but SpinRite still does, and they still work, undocumented, I mean, like deeply undocumented maintenance commands that allow raw data to be written to the drive. And by that I mean that it is possible for me to deliberately write either ECC correctable or deliberately uncorrectable data to a drive. And there is an assembly time switch. It's called the "thrasher." And I'm able to turn it on and specify what percentage of sectors I want this option to destroy. And it's able to go out, and actually it uses a deterministic pseudorandom number generator because I want to, if there's a problem, I want to be able to deterministically recreate the problem.

And so it's able to go out and deliberately create errors that I specify in type and extent and, indirectly, in location, out on the drive. And then I turn SpinRite loose and verify, and in some cases watch carefully, it go through all of its correcting procedures. So I'm able to induce drives to relocate errors that frighten it by how bad they are, but they're not so bad that it can't correct them, and it has to rely on SpinRite. So all of those things that you hear me talk about, I'm able to do reproducibly with perfect, I would say, granularity, or specificity and precision is the word I'm looking for, on contemporary drives. And that's part of what goes into the development of SpinRite. So cool question, Quinton, that I don't think anyone has asked before. I think I've - I know that I've spoken of this, but probably over in the GRC's SpinRite R&D newsgroup. So thank you.

Leo: Time for questions.

Steve: And this is who I'm going to tell the people who write to me, saying, hey, your eCommerce system...

Leo: Can I have your eCommerce system?

Steve: Yeah. Can we please use that? Uh, no.

Leo: That's one thing you don't want to get into, I can tell.

Steve: I'm not doing that.

Leo: Question numero uno comes from Marvin R. He's on the Twitter, @marv51. And he says, have you figured out a way to blacklist sites in uBlock Origin? I mean show ads by default and then turn on permanently when needed? I'm not sure...

Steve: Okay, so...

Leo: What is he talking about?

Steve: Yeah. I selected this because I got a lot of this in the last week. So this is for all of our listeners who said, hey. Well, okay. So first of all, let me explain what he's talking about. The way uBlock defaults is to block a lot, and then you selectively unblock. So the idea is that the user can manually whitelist sites where they want to disable uBlock. And this is true both for uBlock and uBlock Origin.

So what Marvin and many other people wanted was - and, for example, paraphrasing from what I read from other people, they're like, hey, you know, I want to support website advertising. Thanks to all the dialogue between you and Leo, I get it that this is like you're being a good citizen of the Internet if you allow ads to appear on web pages because that's generating revenue for those sites. I want to do that. But if there's a site with particularly obnoxious ads, then that I want to blacklist, rather than whitelisting the ones that I want to allow. So the notion is flipping over the sense of blocking, which is the default.

Now, I did some digging around, and I had to chuckle because this was my characterization, everyone will remember, of Raymond Hill, who's the father of all of these uBlock products, when I described him as sort of what you got if you mixed John Dvorak and...

Leo: Give me a hint.

Steve: Stallman.

Leo: Richard Stallman.

Steve: Richard Stallman.

Leo: Oh, boy.

Steve: If you mix John Dvorak and Richard Stallman, this is kind of what you get. So, for

example, naturally, many people over the course of the last year about, is how long uBlock has been around, many people have wanted this. And so, for example, I found a posting in January that Raymond Hill responded to. In January this guy posted, "Hi. First of all, thanks for the work on uBlock. It works great. I would like to propose a feature where nothing is blocked by default except on the sites where I explicitly enable uBlock. Mostly I don't mind ads, and I want sites to be able to generate revenue through them, except for a few exceptions where they're really obnoxious. I would really like if uBlock supported this mode, which is in essence an exact reversal of the way it works now - on by default, can be disabled forever on a site with a single click."

So that same day Gorhill, which is, we know, the handle that Raymond uses for himself, replies: "This has been" - and think Dvorak and Richard Stallman. "This has been requested before." And then he gives a #177, because I this was like 522 was the posting of this question. So this has been requested before.

Leo: Let me guess. Fuggedaboutit.

Steve: He says, well, but I love this. He says: "Sorry, but I can't personally support this. It's like asking me to give users the ability to opt out of corporate stalking, of being data-mined by default. uBlock serves users' interests, not data miners' interests. Ads," he writes, "are only the small visible portion of corporate stalking. Users should care even more about the portion that is not really visible. Turn on 'I am an advanced user' mode and select only EasyList as the active filter list, and see for yourself all the remote connections, even without ads on the page."

Now, that was his position in January. And I kept digging around, and I think he just sort of finally said, okay, fine. Now, what we really want, we're not going to get. And what we really want is a setting that would just invert the sense of that big power button. Because, you know, the big power button for either non-advanced or expert users is the way we disable uBlock for a given site. So what we'd really like is for it to be off, to have the option for it to be off for sites we have never visited, and if we wanted to, to then click it to turn it on. Which of course is not the way it works.

There is a way to do this. There is a way to give Marvin and everybody else listening who said I want to use uBlock, but I want to flip the sense so that, rather than me whitelisting, I am blacklisting. And I made that the bit.ly link of the week. So you can go to the page, bit.ly - you might want to put this up on the screen, Leo - bit.ly/sn-524, which is today's episode number, 524, bit.ly/sn-524. And if you scroll down, you'll see there are two things that - there's one thing you do to sort of do this globally. And the very first line - oh, I'm sorry, you have to be in expert mode, also.

So our users know that you go to Settings, then select Expert Mode, and then you have to hit the plus sign on the dropdown that you get when you press the button on the toolbar. That opens that extra panel to the left. So in order to do the default Allow All, in the first column, the first colored column of color things, the very top item is All. And so this is something which Raymond put in for this reason, because it's the only reason you would do this. You click it to green. And that's global.

So what this does is it puts in a global override which does come before the static filters. These are called "dynamic filters," the things you do on this panel. And they override the static filters which are the list-based filters. So you would click that first column in the upper line to All. So now you'll get a big dark green highlight and then a faded green for everything else down the column. And what that means is now everything is permitted.

Then, on a site-by-site basis, you do the - on that same page there's like a second example where what you do then is you click the second column on the All also, making it gray, that dark gray, because the first column is global, like global override for all items of, no matter what they are, where they're appearing. The second column is per site. So you click the All in the second column to dark gray, where you want to fall back to uBlock's standard operation.

Again, you know, this is not the way he wants it. It's a little awkward. Frankly, it's much more difficult for me to explain it than it is for you to go to that page. So bit.ly/sn-524. That'll take you to the uBlock wiki. The page is dynamic filtering, turn off uBlock everywhere except, which is the name. You could probably Google "dynamic filtering turn off uBlock everywhere except" in order to find that page, also. And that allows you to do it.

Leo: Why does he block third-party frames? You could leave that in green, too; right?

Steve: Well, yes. And he's done that just because that's sort of what he recommends. He doesn't like the idea that third parties are allowed. But when I tried it, in fact we talked about this last time with you, if you do that on your site, Leo, it just shuts it down because so much is coming from third parties. And I tried it, and it was just like, it just destroyed the web.

Leo: Oh, wow.

Steve: But again, you combine John Dvorak and Richard Stallman. And, you know...

Leo: Let's see, let's...

Steve: You're basically going to get a text-based browser.

Leo: Requests to this server have been blocked by an extension. So you can't watch live. All right. I'm not getting any calendar stuff because that comes from Google.

Steve: Oh, yeah. It's not...

Leo: It's not too bad, though. I mean, you can get - most of the stuff still works. You could subscribe to shows and all of that stuff.

Steve: Yeah. Anyway, so for me - and I have to say this has been just a grand slam win. You immediately took to it two weeks ago when we talked about it.

Leo: Yeah, I like it.

Steve: I've heard a lot from people who are, like, this is really working for them. So that's the story. There is a way.

Leo: Now, how do I get back - I turned all this crap on. Now...

Steve: Click, click the middle on the All.

Leo: On both? Okay.

Steve: Yeah, click, click the middle on both.

Leo: Okay.

Steve: And they should go away.

Leo: Okay. Good. Whew. All right. I want blocking. You know, you say you don't like tracking, but I think there are a lot of people just don't want to see belly fat ads. I think there's also a significant portion of people who are worried about malware being served with these ad servers.

Steve: Yes, yes.

Leo: So it's not just tracking. I mean...

Steve: You're right, yeah.

Leo: ...there are other reasons people use these.

Steve: And, you know, you mention belly fat. Why am I seeing belly fat ads? I'm seeing them, too. It's like, and there are...

Leo: You need to turn your tracking on so you don't see belly fat ads. That's the point.

Steve: They're obnoxious.

Leo: See, people, I - well, I don't want to start. I don't want to start. Ashton C., @ashtonc on the Twitter: Steve, I thought you might be interested in the latest ad network tactics to get around adblockers. Those evil advertisers. A number of my clients have been asked by numerous ad networks to add customer reverse proxy configs to their web servers, like Nginx. These configs essentially say, if I can't find any image file locally, instead of saying 404, proxy the request to the ad network's servers and return the ad network's image. This makes it look like the ad image file is coming from the original site instead of a third-party ad network site. I feel that this practice is very sketchy. I've declined on all such requests. But of course this is just the beginning, if people start really using these blockers.

Steve: Yes. And, I mean, we're going through, there's no doubt about it, we are entering a fascinating time. One of the things that didn't make my notes for this week's news just occurred. It was in the news today. And that's Apple and the federal judiciary. I don't remember which branch or block. But Apple is now fighting a subpoena for iMessage messages and turning down the federal government's requests, saying we're sorry, we can't give them to you. We cannot comply with your subpoena because iMessage is encrypted, and we don't have the keys. And Microsoft is similarly fighting a request. In this case it's a request for email from servers in Ireland. And Microsoft is saying, uh, they're in Ireland. They're not here. And we contest your jurisdiction. And the privacy rights over in Europe, their privacy laws don't allow us to provide that.

So again, on all fronts on this whole issue of the encryption of user content and data, the federal law enforcement wanting backdoors installed in order so that Apple could be compelled under subpoena to produce output of iMessage conversations, and then the whole question of the viability of advertising and how the industry's going to evolve and handle it, all just fascinating, you know, pieces of technology that we're perfectly positioned to explore.

So this is something, I got a kick out of this because I hadn't been aware of this happening, but I've been talking about it happening. Because from a technical standpoint this is clearly a means for advertisers to circumvent adblockers. And that is exactly - here's an example of it being done, which is, if the browser asks for content which is unavailable from the first-party site, the ad networks, Ashton says, are already asking sites to configure a proxy so that the server will turn around and ask the ad server for the content on its behalf. It receives it and then returns it to the browser. So the browser is making image requests to the first-party site, which it's forwarding, which that first-party server is forwarding to the ad server and returning them.

Now, frankly, for images, this is not a problem. I mean, this is something I can expect to see. And I noted that Ashton did say "images." However, what we're seeing is script. And this would be deadly if this was done with script because what would then happen is that it completely breaks the same-origin policy of browsers. One of the things, one of the most important protections we have in our web browsing, and we've talked about this often, is the same-origin policy, which means that, if sites do retrieve script from third parties, and we know they do, I mean, you know, all of this gunk that's coming in from third parties is typically not an image. It's script which wants to generate a supercookie to strengthen tracking beyond a third-party cookie, or it wants to do the ad serving, rotating stuff itself, rather than pushing that off to the ad server.

So what's super critical is that browsers obey the same-origin policy, meaning that script coming from another domain is restricted to only querying from that same other domain, that is, what it's basically doing, it's stovepiping all of these domains, in some cases 48

different domains are dumping crap into our browser. Well, okay. But do not allow them to talk to other servers, and especially not to the hosting server because its script needs to know that it's not - because, I mean, this just opens up a huge possibility for malicious script.

So proxying like this we may be seeing. But I sure hope it never happens with anything but an image. I mean, even an image is scary because we have seen exploits using JPEG rendering mistakes that allow a remote code execution, essentially an image contains executable content which, due to a fault in the way the image is rendered, allows that to execute. So images can be a problem. But they're way less a problem than scripts. And so the idea that a first-party site would host script which would come from a third party is just blood chilling. So I expect we're going to see unblockable ads moving forward. I just hope we don't see unblockable script because that really does change the game.

Leo: Now we continue with another question for Steve. This is from Evan Drosky, Upstate New York. He's defending Lenovo, well, at least a little bit: Steve and Leo, in Episode 522 you mentioned the Lenovo BIOS fix that they released, and you both said Lenovo is basically dead to you. However, I work in an IT department that uses certain Lenovo computers, and I want to make the distinction that the ThinkPad brand of professional computers that they bought from IBM and continue to produce were not affected by the BIOS file resurrection problem nor the Superfish fiasco.

We've used ThinkPads and ThinkStations since before the Lenovo purchase of IBM's desktop computer business, and I can say that their stewardship of the well-regarded Think brand has been and has remained well above par. I would actually agree with you on that.

Steve: Yup.

Leo: These systems have always been dependable for us when treated right, and never been prone to the kind of problems Lenovo's been putting into their in-house lines of consumer products. Other than the standard bundled software like Office and trial versions of antivirus, the level of bloatware on these systems is less than the admittedly horrible industry standard for bloatware.

Yes, when left to their own devices they've soiled their reputation with these issues on their consumer line, and I agree 100% with staying away from them. But the Think line of professional computers have always been a separate issue, and thankfully they seem to be keeping these things separate. I wanted to pass this along before everyone completely writes Lenovo off. Hopefully - fingers crossed - this kind of stupidity will never reach the Think line of computers. Thank you both for the show, and here's to many more years to come!

You know, I had heard, and I haven't confirmed this, that IBM on sale of ThinkPad to Lenovo insisted on some oversight of the brand for some years.

Steve: Interesting. Interesting. I know...

Leo: I've also heard that that oversight has expired.

Steve: Oh. Well, you and I have both mentioned that, as far as we knew, the ThinkPads were okay. And they are, you know, my two Lenovos are ThinkPads, and I love them. And so I just wanted to make sure, in the interest of keeping the record straight, that we didn't just with too broad a brush wipe Lenovo out because I think - although there are now some good alternatives, too. Dell is producing some nice...

Leo: There's plenty of good alternatives, yeah.

Steve: Yeah.

Leo: But there's some - the ThinkPad X1 Carbon, I mean, they're...

Steve: They're, I know, oh. It's...

Leo: They are lustworthy. I mean, there's some really nice...

Steve: They are just so beautiful.

Leo: Yeah. And I think the question, you know, just keep your eye peeled, and we'll find out if stuff like this has happened, I think.

Steve: I would say yeah because, I mean, go in with your eyes open.

Leo: Yeah. Question 4 from Niall in the U.K. He's been - something's been haunting Steve: What is the true value of tracking?

Steve: Hmm.

Leo: Oh, interesting. I've enjoyed the back-and-forth between yourself and Leo on advertising. I can see both sides, and I'm not adverse to the advertisement model, although I have never ever knowingly clicked on an advert on a web page. But what I abhor is the tracking element. As you said, Steve. I have searched for references that detail and justify the effectiveness of tracking as a revenue booster, but I can't find anything definitive. Given the huge efforts to track me and the ecosystem that has infiltrated the web, it must make a significant impact on the ability to serve me relevant web ads and subsequently make me buy more stuff. Right? Or is it all smoke, mirrors, and sales talk?

Steve: And you know, I love this question because, as I said at the top, it's been haunting me. I'm going to try, I'm going to see if I can find somebody who actually knows and ask him to join us on this podcast. Because I think you and I and our listeners would really benefit from hearing from someone who really does understand what this is. And...

Leo: The trick is to find somebody who's going to be truthful about it.

Steve: Yeah, I know. I think I may have a line on some ex-Googlers, and they might be in a position to know and be unbiased. But, I mean, again, we would know who they were.

Leo: I've heard research, and I've seen evidence that tracking isn't effective. It's not that it's - but that's kind of moot because, if you're a website that has ads, that sells ads, the advertisers insist on at least enough tracking to count the impressions. That's how you pay them.

Steve: Yes. Although that's not...

Leo: And by the way, you can't tell what information they're tracking, really; right?

Steve: Well, see, but, okay. So there's counting, and there's tracking. Counting, no one has any objection to. But tracking means cross-site, and that's what creeps people out is this notion that, you know, because many different sites pull ads from the same source, and the nature of the way the web works with cookies and supercookies and all that, that it's the trans-website which is tracking. Not counting. Counting is fine. Counting no one cares about.

Leo: But you can't distinguish between tracking and counting as an end-user. You don't know if that tracker is doing more than counting because what they're doing is they're capturing your IP address. So at that point...

Steve: Well, they're...

Leo: And they need that, by the way, to do unique tracking. They also, I mean, I'll tell you a few benefits. For one thing, we know that, advertisers know that there is an ideal, I've mentioned this before, an ideal number of impressions. You don't want to make too few, but you don't want to make too many, either. So one of the things they do with counting is also make sure you don't see that belly fat ad too many times. I think some of the problem is the choice of the word "tracking."

Steve: Well, yes.

Leo: Which implies somebody's kind of snooping on you. When in fact what they want to do is find out what your interests are so they can serve you ads for something you'd buy.

Steve: Okay. So first of all, they're not using IP because, for example, everybody at Starbucks has the same IP. Everyone in a corporation behind a NAT router...

Leo: Okay. So they use something, some supercookie or just a plain old cookie to identify you.

Steve: Yeah.

Leo: But there's some unique identifier, yeah.

Steve: Yes, exactly. It's just a cookie. And in listening to you and me talk about this, I've sort of thought about this. And what I wanted to do was just to say that, you know, you are of the opinion that people are getting way too worked up and upset and concerned about this. And my feeling is that, you know, I'm not that concerned about it. But my sense is that advertisers are tracking us because they can. And users are saying "no thank you" because we can. And so it's not, you know, getting worked up into a big froth. It's just that, hey, if it's easy for me to deny them the ability to track me, and I don't see any great value from it, then, yeah, I'd rather not be tracked. So not a big deal.

Leo: Yeah, I can understand that. I don't - there's two things I would say. One is I hear a lot of kind of generalizations about why tracking's bad. I mean, I understand, if you don't want to be, you shouldn't be. But I hear generalizations. I have yet to hear a really serious issue with tracking. Like people say, oh, your insurance company will track you to Dunkin' Donuts and then - they don't need to do that. Trust me, they know more about you than any advertiser can give them, partially because you had to give them that information before you got insurance; and if you lied, well, they'll deny you payment when you need it. So that's not - I think that that's a - there's a lot of straw man arguments about why tracking is bad. I'd like to hear a real argument about the detriment of it.

But secondarily, and more importantly to me - and, by the way, if I were arguing in my own economic self-interest, I would say block, block, block because the fewer, the worse ad banners work, the more ad revenue is going to come to me. So I think, you know, everybody should block as much as possible. But what I worry about is the - here's a guy, the How To Geek. He started writing great little, useful little articles about, you know, useful tips on how to use computers and technology. After a while he put Google AdWords on there, and he was getting a few hundred bucks a month. He said, this is nice. He put a few more ads on there, realized he could make enough money with these advertisements to quit his job and really go all in on this website.

Steve: Do a much better job.

Leo: There are 50,000, 100,000 websites just like this. As soon as tracker people - the key to this is, if you're running a blocker, you're now invisible to that site. You go on that site, and that site has no way to know you visited, basically, because you've turned off number counting. You've turned off tracking. Most people, if you're using Adblock or uBlock Origin in its default, you're invisible. You go on that site as a ghost. If 20% of his readership does that, that's a direct - forget 21 billion.

Steve: Yup.

Leo: That's a true - and I don't know what the number is, but that's a direct 20% loss to him in revenue, period. If 50% do it, there's a 50% loss to him in revenue, period. It looks to him like half as many people are visiting his site, is what it looks like. He has no way of knowing if that's what that number is, or why that. Maybe twice as many people are visiting, but they're using blockers. He doesn't know because you're invisible now.

Steve: Well, he does know because he knows you've visited. And so, I mean, so he knows what his traffic is. And so you're right, what he will see is, he will see a dropping in revenue from the people whose ads he's hosting, yet maybe even an increase in his actual site traffic because, you know, he knows what his site traffic is because he's still serving his pages.

Leo: If he looks at his logs.

Steve: Yes.

Leo: If he looks at his logs, he knows, yeah. A lot of people don't know how to look at their logs or can't see their logs. But most, if you look, one of the reasons you see Google Analytics on almost every site you visit is that's how he's measuring his traffic. He's not...

Steve: Right, right.

Leo: Google bought, you know, the Webalizers of the world, the little utilities that you'd run against your log. I used to do that to find out my - but Analytics is so much more useful, most people stopped doing that. Maybe they'll go back to doing it.

Steve: I see. So you're right. So you're right. If the adblockers are blocking Google Analytics, and it is on the list, then you're right, you disappear.

Leo: So you might have to go back to parsing your blog, which is a nontrivial thing to do. And I bet you most, I mean, just look at the sites that use GA. We use GA. Everybody uses GA.

Steve: Yeah.

Leo: And GA bought all the little utilities, Urchin and Webalizer, and all of those are owned by the big companies.

Steve: So your second argument I completely agree with. And that is, I mean, the actual damage being done by web blocking. Yeah, by adblocking of tracking. The first issue I don't - and that was what I was aiming at, was that you're looking for, like, some big harm that's being done. And I'm just saying there doesn't need to be big harm done. I just don't want tracking. I don't want to be tracked. I don't mind being counted. Being counted for advertising is one thing. But having a cross-site tracking profile, where I go to a different site, and the ads know who I am, ostensibly because they've got a profile on me, that I'd rather not have.

And so, again, it doesn't have to be the end of the world, where insurance companies are going to deny us coverage. It's just I don't want to be tracked. And that's a preference of mine. And it's easy to make that happen. What we really need is to be anonymous. And of course the other point that you've made is, when you're on the Internet, it is very difficult to be anonymous.

Leo: Yeah, I think you're not really anonymous.

Steve: Truly untrackable.

Leo: Come on. But anyway, I mean, that's like saying I'd like to walk through the mall and have no one know I'm there. I mean, you can cloak yourself, but people are going to [indiscernible]. Anyway, that's my point is that I feel - it doesn't affect me. In fact, it would be beneficial to me if everybody blocked ads, if all of those websites just collapsed, because you'd have to turn to us and our ads. But in principle I don't want people to start blocking ads because they can block our ads even easier. You just fast-forward through them.

Steve: Yeah. We're going through an interesting time.

Leo: And people are pissed off at me for bringing it up. So that's the last time I'm going to talk about it. Alan Farough in Kanata, Ontario, Canada wonders how data is stored on glass platters. Well, we just talked about that at the beginning: I've been dimly aware that spinning hard drives use platters made from either metal or glass. But the question that had never occurred to me before is how is data written onto a drive with glass platters? Never heard this question discussed before. I figured if anybody would know, you would. Explain, Mr. Gibson.

Steve: So historically platters were made of aluminum. And it turns out that just by virtue of the difference of the characteristics of the material, as densities increased and head flying heights lowered, it became increasingly important for these surfaces to be unbelievably smooth. And it is possible to manufacture smoother surfaces with glass than with aluminum. There's just the manufacturing process, who knows what it is, the annealing or whatever.

Of course glass, as we know, is actually a fluid. So that, for example, if you go to very old churches, sometimes the windows at the bottom of the sill are thicker. The glass of the windows is thicker at the bottom than it is at the top because over hundreds of years the glass has flowed because it is actually a very slow-flowing fluid. So as a consequence of this, manufacturers who are looking for an alternative to aluminum turn to glass. Glass

is stiffer per unit weight. And that allows them to make platters which are thinner for the same strength. And of course thinner platters allow them to put more platters in the same Z axis, in the same height of the drive.

Now, there is a problem with brittleness because, rather than giving nicely, as the picture at the top of the podcast shows, glass notoriously fractures. And so now what's being done is manufacturers are experimenting with glass-ceramic hybrids in order to make them stronger. I mean, there's still going to be a breaking problem.

But the good news is many people are asking, how do I absolutely destroy this drive? One thing you can do that we've never mentioned is the same way some people prepare taffy before they eat it is you put it in your hand, and you slam it down as hard as you can on cement, which is a harder surface even than asphalt. And some drives will absolutely shatter. And if you then shake the drive, and it sounds like a million little bits, then what you have is glass platters, and you've just pulverized them without having to run DBAN or magnetize them or drill through them or do anything else. Just giving them the old taffy fracturing smack on a hard surface. And if it powders the disks, then your job there is done. Oh, and as to how data is stored on glass...

Leo: Yes, wasn't that the - oh, sorry, I didn't - okay. Go ahead.

Steve: Yeah . Neither aluminum nor glass were ever the data storage medium.

Leo: It's just a substrate.

Steve: They were the substrate.

Leo: Yeah.

Steve: Yes. They were the substrate, just the mechanical underpinning. It used to be that we would deposit oxide, ferrous oxide of some sort, often iron oxide in the old days. And the deposition materials have evolved. Now there's this amazing multilayer sputter deposition technology. They're actually putting several layers on, then the ferrous layer, which stores the data. And on new drives they then put a lubricant layer on top of that. So, I mean, it is a multilayered process.

So glass itself, as Alan wondered, could not store magnetic material, and it never did. Neither did aluminum. The aluminum, just like the glass, was only the underlayment used to hold the actual storage layer, which is a microscopically thin layer of something of a magnetic medium, so it doesn't matter whether it's aluminum underneath it or glass underneath it. It's actually deposited itself in a superfine smooth coating. What you want is no lumps and bumps underneath. And glass, it turns out, they're able to make much smoother even than the smoothest milling of aluminum.

Leo: Rob Best, Colton, New York wonders about proxies, caching, and advertising. Oh, god. Here we go again.

Steve: This'll be easy.

Leo: I'm shutting up. Steve and Leo, I've been following your interesting discussions about the ethics and mechanics of adblocking. And thanks to those, I've been changing my own behavior to better support the sites and services I use. That's nice. Thank you,

Steve: See, I really do think that people, like when we talked about Ars Technica, I mean, there is a goodhearted willingness and a desire to support the web. In other words, we just don't want to be hurt as a result.

Leo: Yeah, no. And I, you know what, I support that. And frankly, I'm hooked on uBlock Origin.

Steve: Yeah.

Leo: Thanks. What I'm wondering is how do proxies and caching servers impact content and service providers? Advertisers want accurate location and event-count information; but won't proxy and cache servers give them false or misleading information and, in the case of caching servers, give low numbers of downloads? For instance, the first client behind a cache server counts as a hit, but subsequent clients just get the content from the cache server, not from the true source. As always, a happy listener and SpinRite owner. I had the school I work at buy a site license a few years back.

Steve: So, interesting question. And there's two answers. First of all, the web has long ago dealt with this problem because, on one hand, you want caching when you want it. Like for a website's static images, like the decoration of the page, the site's icons that never change. The site doesn't. It's not in its interest to be resending those over and over and over, if, for example, the user's ISP has a caching server. And we know that many ISPs do have a caching server.

I know that Cox, for example, my cable modem supplier, they have a caching server which I had to bypass in order to use - so that ShieldsUP! would see the connection coming from me, not from the server, because otherwise they were trying to filter requests for web pages, which might be static. So not only images of static pages, but the pages themselves.

So all assets which are served by a web server have an expiration date. And that allows the web server that is serving this content to control how long a cache anywhere in between the server and the user's eyeball is allowed to cache the content. Even browsers have a cache in them. So not just external proxies and caching servers, but your browser itself. I mean, that's that image cache that has gotten people in trouble in lots of interesting made-for-TV movies and so forth, where somebody checks the browser cache to see what Junior's been doing because the browser is saying, hey, as long as this image hasn't expired, I'm going to hold onto it in case the user goes to a page that asks for it again, because then it's instantaneous. Don't even have to send a request out.

So it's with the expiration date, this so-called metadata that goes along with all web

assets. They have an expiration date. And if you explicitly wanted to count something, you can set the expiration date into the past. That's an officially sanctioned thing to do, to set it to 1970. And so here we are in 2015. Something that expired in 1970? Well, when it's fresh, it's old. So that's going to cause the web browser to ask for that again and again and again. There are also other meta headers which can be put in, specifically that says no caching. So the browser and anything along the way can be explicitly told, independent of what day it is or the date, just do not cache this. I don't want it cached.

And then what we're seeing, and this is one of the tricks that blockers are having, is serial-numbered assets. That is, even with all of these other mechanisms around, what we're beginning to see, because we've just got excess computing power and technology oozing out of us, is that servers are issuing images with changing numbers. And that's the ultimate defeat of any sort of cache or proxy because every single time they serve something, they generate what looks like a pseudorandom number. I mean, it's long, and it's just gibberish. And you refresh the page, and all of those change, which means they all have to be fetched again because, as far as the browser knows, they're all different items. So this whole problem of any kind of proxy or caching getting in between the user and the browser, that's been, you know, that was part of the original plumbing of the Internet. We've gone further than the original plumbing. But the problem's long since been solved.

Leo: Question 7 comes to us from Gregg Nicholas in Michigan, United States of America. He says: I still don't trust anyone. I've got to disagree, Steve, with your position on allowing vendors to autonomously make changes to my equipment. Several vendors have shown that if they can "patch" equipment without our permission, they can remove features or functionality and add spying features. If we allow vendors to modify our equipment when they choose, then we no longer really own it. The equipment doesn't serve our needs, and it will be modified as they choose to increase their profit.

Steve: Okay. And Gregg is actually Richard Stallman.

Leo: Well, when did you even say that? I don't know, do you have that position?

Steve: Yes. And I still hold to it. It was one of my favorite features of the Google OnHub was that it would be dynamically reaching out and updating its firmware.

Leo: Oh. I do remember you saying that, yes.

Steve: Yeah. And I really - and in fact I even asked our listeners to come up with an acronym. I got a lot of submissions, that is, an acronym like TNO which would be pithy and fun, and it would convey the idea that a device can be updated on the fly. That is, what we keep seeing, for example, is routers with firmware problems, or phones, the Android phones we've been talking about where they've got problems, but they've been abandoned by their suppliers because they're a couple years older, and they don't want to bother with keeping them current.

So my feeling is that, try as we might, we are unable to, I mean, history shows us we're unable to create secure solutions. So if it's going to be on the Internet, the person who

created it needs to retain responsibility, and the device needs to be able to be responsible for itself and update itself. Gregg is taking the, I mean, a reasonable alternative position, which is I don't want my stuff talking behind my back and being updated behind my back. I see both sides.

And I really think we're talking about, again, this is my "we're not all the same." If Gregg is willing to assume responsibility for the management of his stuff, then I think he should be able to say, no, don't self-manage. I want to take responsibility. But most people are just going to say, I don't want to be updating my light bulbs. Let them update themselves, if they are able to. Because if a problem is found with my light bulb, I'd like it fixed as quickly as possible, thank you very much.

So there are, I'm convinced there are secure ways of allowing our Internet of Things, the IOT devices, to keep themselves current. They can ping to see if there's a new version of their firmware and arrange to do this securely. And largely I think that's probably a better idea. Yes, it creates a vulnerability because you could ping and get a malicious download. But again, there are ways to do that securely. I just think that's where we're headed. And on balance, we're probably better off.

Leo: Unless you are writing your own source code and compiling it yourself, there's a certain degree of trust. I mean, if you're using any form of closed source software, you're 100% trusting the vendor.

Steve: Yes. If you're using BitLocker on Windows to do your encryption, you're assuming Microsoft hasn't built in a way to reply to law enforcement.

Leo: And I'll go a step farther. I mean, living in society requires trust of your fellow man. As you drive down the street, you trust that the person coming the other direction is not going to swerve into your lane and clobber you. And if you didn't have that trust, it would be very hard to drive. So that's called being in society. And if you really want to be safe and secure, you need to remove yourself from society, write all your own code. Oh, by the way, you might want to write the microcode in your processor. And what about the BIOS firmware? Stephen Brooks...

Steve: And my...

Leo: Go ahead.

Steve: Yeah. My only point is there are also sort of compromises, where you just, for example, with Windows 10, you turn off all of that junk. I mean, how many apps? There's three more apps this week for, like, making Windows 10 stop spying on people. So obviously there's a market.

Leo: I hope you trust those apps. Did you compile them yourself? Did you examine the source code? Stephen Brooks in Red Bluff, California wonders how to simultaneously be both a Google Contributor and a uBlock Origin user? I signed up for Google Contributor - that's that system we talked about where you pay Google

some bucks, and then they feed them out to the sites you visit, in lieu of running ads on those sites. I signed up, but how can I whitelist for Google Contributor advertisements while using uBlock Origin? Enjoy Security Now!, and a proud owner of SpinRite.

Steve: So I sort of put this question out to everyone. I don't have uBlock on my iPad yet. And when I'm out and about, I see these, I mean, I'm always seeing what I now recognize as my Google Contributor ad replacements. iMore is, like, that's what they use. I get these really nice pastel circle things all over the iMore.com pages. When I'm home, I don't see them because I've got uBlock. And just like Stephen, I would love to allow Contributor to still work. And this is an example of the kind of solution I think we're going to find. We're in interesting times, as they say. We're going through an upheaval. Adblockers really are having an effect because users really are beginning to adopt them, and they're becoming easier to use.

So I don't have an answer to this. But sign me up for allowing Google Contributor to poke a hole through uBlock Origin. And if any of our higher end tinkerer listeners figure out how to do that, make sure you bring it to my attention, and I will share it with everybody because that would be neat. I really like Google Contributor. I love supporting the sites I visit. In fact, Elaine loved the W3Schools site that I mentioned was another place where I see the same Google Contributor things because I was doing JavaScript for SQRL about a month ago, and she hadn't ever run across W3Schools and really likes the site.

Leo: Christian Steinway, Dallas, Texas has a thought about Lenovo and its pre-boot tricks: Wouldn't the order of events at boot-time, with an encrypted boot partition, prevent such a reach-up into the file system? Isn't that kind of what UEFI and the whole, you know, trusted boot and all of that, secure boot, aren't they to protect that kind of tampering?

Steve: So what Christian is talking about relative to Lenovo is this creepy thing that we discovered they were doing, which was they were, if your drive was a FAT or an NTFS file system, it was able to parse the file system and replace a Microsoft file in the Windows\system32 directory. And he is completely right. And I just - I saw this, and I thought, yeah, I should have mentioned it. So I'm mentioning it. If you have encrypted your boot partition, that won't work.

Now, the bad news is that next-generation thing that Microsoft has officially sanctioned, where the firmware can write the EXE up into RAM and then, through the ACPI table, tell Windows that there's an EXE that it would like run on its behalf. That will still work because that happens after the boot partition has been decrypted, and Windows is running. So the creepy direct manipulation of the file system, exactly as Christian suggests, that's defeated if the boot partition is encrypted. But unfortunately, what Lenovo was most recently doing, and the technology for BIOSes to do this, which Microsoft does sanction, unfortunately cannot be bypassed with a boot partition encryption.

Leo: Google with Chrome OS does this, I don't know what, a checksum, a hash or something on the OS. And it validates the OS before it boots that it has been unmodified.

Steve: Right.

Leo: Would that be effective in a case like this? As long as you kept stuff from running before that?

Steve: It would depend upon what they're doing and when because, for example, the secure boot does do a good cryptographic hash signature of what's there before it runs it. Unfortunately, this is replacing a file. And so Windows is not checking the entire file system that it hasn't been changed. It's saying, is this component that I'm about to load into RAM, that I'm reading from the file system, has it been changed? So basically it's a digital signature per file. In this case, it was a file that was still run, even if it was not signed.

Leo: Right, right. Finally, our last question comes from Thilo Maier in New York City, who wonders about plugin precedence when using uBlock Origin and Privacy Badger at the same time: Steve, ever since you mentioned the website PrivacyTools.io, I have been testing uBlock Origin. After acquiring a taste for it, I like it more and more. Yes, I do, too. At the same time, I like the idea behind the EFF's Privacy Badger, and so I have both extensions installed. What I'm wondering is, is that a good idea, and whether or not the two similar extensions might interfere with one another.

For instance, if I visit TheVerge.com, uBlock Origin shows domain amazon-adsystem.com as blocked, but Privacy Badger doesn't recognize it as a tracking domain. Can I be sure that amazon-adsystem.com is blocked if one extension blocks it? Or is there such a thing as extension precedence - I think you're overthinking this - in which one of the two extensions will be executed last and will win? Two enter, but one comes out.

Steve: So just to remind people, the different behavior is that uBlock Origin has a curated list of domains that they have decided they want to block. Privacy Badger is behavior-based blocking, where based on their description of Privacy Badger, if Privacy Badger sees a domain tracking on three different occasions on different first-party websites, then it sees the tracking behavior and starts to block that particular third-party site. So they operate differently and provide different functionality just because they're sort of taking a different approach.

The way extensions happen, and this sort of comes back to that bug in Chrome, where Chrome was supposed to be notifying extensions before loading a given asset, that happens in a, sort of as he suggests here, in some precedence where each extension that has registered for the notification gets an opportunity to say, no, block that. So it's not a majority voting process. It's an every single extension has to say yes. It's like old-fashioned light bulbs on the Christmas tree. If one light bulb is burned out, the whole strand is dark. Similarly, every single plugin that has said "I want notification before loading of assets," every single one of them has to say, yeah, that's fine. Yeah, that's fine. Yeah, that's fine. Yeah, that's fine.

And then so Chrome is continually sending that notification out to one plugin after another. And if all of them say they have no problem with it, then Chrome says, whew, and then sends the request out to wherever. So the answer is multiple plugins have an "and" relationship with, you know, from a logical "and." They all have to say yes in order

for the browser to load the asset. If any one of them denies it, then the browser will not load. So that's how precedence works.

Leo: Nice.

Steve: And you can play with it yourself, just by disabling Adblocker, or disabling uBlock Origin and then watching that Privacy Badger will block the site if you can, like, bounce around awhile and show ads that cause it to then turn on blocking. Then it'll block it. Then if you turn on uBlock Origin, then you uBlock Origin may or may not still show it blocked. Because if it's blocked before it gets to uBlock Origin, sort of in the hierarchy of sort of this chain of plugins, then Privacy Badger will have blocked it upstream of uBlock Origin. But everybody has to agree it's okay in order for you to see it.

Leo: And Privacy Badger also has a larger political goal, to enforce, make Do Not Track enforceable and encourage people to support Do Not Track, et cetera, et cetera.

Steve: Which, you know, maybe that would be a fabulous outcome.

Leo: Yeah.

Steve: If we could just turn on DNT, and the ad companies, because users flexed our muscles...

Leo: Yeah.

Steve: ...and said we don't want you tracking, so you have to stop tracking if we tell you not to track. But then you're still welcome to serve us ads and count us.

Leo: Make enforceable. Now, here's a question. How many of you, if that were the case, you could turn a button and Do Not Track would be turned on in your browser, as it is in Safari and other browsers - not Chrome, obviously. But let's say you use Safari. You turn on Do Not Track. And let's say that were honored by all the big websites, every website you go to. Would you then stop adblocking? Or would you like to still adblock because everything's faster, the size of the page is smaller, there's no Flash ware or malware? I think people would still block ads.

Steve: I think we have a problem.

Leo: I think we...

Steve: Houston, we have a problem.

Leo: Houston, we have a problem. Steve Gibson. This is the guy, man. You've got to go to GRC.com, check it out. Lots of great stuff there, including of course SpinRite, the world's best hard drive recovery and maintenance utility. You could find the podcast there, as well, 16Kb audio, full-quality audio, transcripts written by a human hand, Elaine Farris, who does such a nice job of those. It's all at GRC.com.

Questions for Steve, lots of ways to do that. One of course is to go to GRC.com/feedback. There's a form there. But you can also tweet him. He's @SGgrc on the Twitter and apparently pays attention to that stuff. We also have high-quality audio and video of the show, if you want to see Steve's smiling face. That is at TWiT.tv/sn. You can watch the show as we do it live about 1:30 p.m. Pacific, 4:30 Eastern, 20:30 UTC, every Tuesday on TWiT.tv. You can be in the studio audience, if you'd like. Hello, studio audience. All you have to do is email tickets at TWiT.tv. We'll make a seat for you. But of course, if you can't be here during the live show, you can always listen after the fact. On-demand audio and video is available, as I said, on Steve's site and ours. For all of our shows. Thanks, Steve. Be here next week for another fabulous edition.

Steve: Of Security Now!.

Leo: You're waiting for me to say that.

Steve: Live long and prosper.

Leo: Of Security Now!.