iOS Content Filters

Description: Steve and Leo cover a busy past week of security news, then discuss the first week of iOS mobile web content filtering made possible by Wednesday's release of iOS v9. They take a close look at the initial set of content blocking apps available for iOS and Safari.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-526.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-526-lq.mp3

SHOW TEASE: It's time for Security Now!. Oh, man, is there a lot to talk about. Steve Gibson is here. We'll of course talk about Audi and VW, and we're going to talk about adblockers in iOS. There's a great show ahead. Stay tuned. Security Now! is next.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 526, recorded Tuesday, September 22nd, 2015: iOS Content Blockers. It's time for Security Now!, the show that protects you and your loved ones online, privacy and all that. Steve Gibson is here. He is the man behind GRC.com, SpinRite, the world's best hard drive utility, and of course an expert on all this stuff because it started by him having to deal with spyware and DDoSes and bad guys. And now I think it's got to be, Steve, something you actually are very interested and love doing.

Steve Gibson: Oh, well, can you say SQRL?

Leo: Yeah.

Steve: Yeah, I mean, it's like, this is, essentially, I got into it because it was a new technology where everything was available. What I really like is, well, very much sort of the whole open approach. With the 'Net we had RFCs that laid out how all of this worked. So I was able to just sit down and read through it and figure it out and then begin to play with it. So I got sucked into it because it was inherently a nonproprietary technology that anyone who was interested could just learn about as much as they wanted to. And then it was things like the logic of the software firewall, the idea of finding something in my computer, the adware that I hadn't installed. And it's like, whoa, wait a minute now, then wanting to monitor that and so forth.

So, yeah, it's just, I think, the whole security thing, I mean, we really started, now more than 10 years ago, before it became the huge issue that it is. We knew it was, you know, I mean, you suggested this because you knew that was sort of my focus, not that it was going to be like this crazy hot-button issue where, well, for example, one of the things we're going to talk about is India's trial balloon they have floated up on their national government-backed encryption policy, which is just, whoa, boy. Well, we'll have to see how it goes.

But we're going to talk, of course, about iOS's XcodeGhost problem; the critical Adobe Flash update; AVG's privacy policy update, and calling it a privacy policy at this point is a bit of a tongue-in-cheek; the Cisco router problem, where malicious firmware has been discovered by I think it was the guys at FireEye; the Ashley Madison password mistake that was discovered. We have to talk, just because we've been dealing with vehicle security stuff, briefly about the VW/Audi catastrophe, where they got caught, we think, cheating. A little update on StageFright. And I've taken a look at a collection of the current adblock offerings, all the ones that appeared immediately and have sort of a feature comparison and some recommendations.

Leo: For iOS specifically, not...

Steve: For iOS, yes. Yeah, thank you, specifically for iOS. So I think a great and interesting podcast.

Leo: Can't wait. I'm excited. In fact, I'm really interested in your recommendation for iOS adblocker. I've been so happy with uBlock Origin. We still think that's probably the best.

Steve: Yup.

Leo: Well, I guess there's many good choices on the desktop.

Steve: Right. And in fact the simplest and easiest to use choice for iOS is by Chris, who is maintaining the non-Origin version of uBlock. So he's the guy that is keeping uBlock going at uBlock.org. And he produced Purify, which is his iOS blocker. And I like it because I think it's very clear that what we're going to see - and in fact you guys were already talking about it on MacBreak Weekly, and we predicted this on our podcast, is sites start becoming increasingly sensitive to their visitors' use of adblockers. And some may say, hey, you know, I noticed you have an adblocker running. Would you consider disabling it because we depend upon your seeing ads in order to support the site. Others might deny you access until you lower your shields.

Leo: And I think that's a big mistake, frankly.

Steve: And as was said on MacBreak Weekly, many people sort of flip their finger and go somewhere else.

Leo: Yeah, bye-bye.

Steve: And one of the things I want to talk about was Dave Winer did an interesting post-adblock blog. We mentioned him briefly a couple weeks ago because he was feeling that suddenly his Mac OS was becoming un-usefully social. He was the one who talked about, you know, your baseball bat reminding you that your mortgage payment was due. And he said, "Who asked for this?" Anyway, so he had some little pithy comments that I want to share. But what was interesting was reading that and having thought about all this brought me to a question I have seen nobody ask anywhere. So we're going to ask a question that I think it'll be something you pick up on, Leo, and that you'll be asking your own co-hosts in future podcasts because that's an interesting issue that no one has touched on before. So, yes, a great podcast.

Leo: Well, as usual...

Steve: Yo, Leo. So the picture of the week I just stuck in here. I checked, I wanted to check the connectivity before we started the podcast. And I got a big kick out of seeing that, when I went to Speedtest.net, the introduction page had the standard little Speedtest gizmo in the center. It was surrounded by what I chose as my ad filler for the Google Contributor.

Leo: Oh, these little bubbles or whatever.

Steve: Yeah, they're sort of pastel. And actually they're mouse-sensitive. If you mouse over them, they sort of rearrange themselves. And I'm just bringing up the itemization, and I thought I would just see what Speedtest showed. Ah. Speedtest.net, 61 cents of my money has gone to support them.

Leo: This is Google Contributor, which is...

Steve: This is Google Contributor.

Leo: If you're using an adblocker, you really ought to consider something like this, as well.

Steve: Yeah. I think it's just the, you know, I give them as much money as they will take per month, which is $10. And they spread it around based on my surfing. And now, when I look at the sites I've contributed to, because there's a complete enumeration of them from most money to least, I mean, I'm scrolling down through, like, everywhere I go. It's just, like, crazy. So...

Leo: Well, I'm sure Adblockalypse has pushed more people in this direction, too. They're trying to figure out some solution to the widespread use of adblockers.

Steve: Well, and we're just going through, as we've said on this podcast for the last month or so every week, a transitional phase. It'll be really interesting to see how it all shakes out.

Leo: Hey, can I recommend a better speed test, though?

Steve: Yeah.

Leo: Because Speedtest.net is notorious for being - you know what Volkswagen did with the diesel engine?

Steve: Yes, special casing. This is the problem. Actually, any popular speed test suffers from this possibility, and that is that your ISP will give specific IPs that are known to be associated with the speed testing servers preferential performance. So the good news is I haven't mentioned this on the podcast, but now when I upload the edited audio for Elaine, I trigger GRC's DOS attack alarms because my connection is so fast to GRC that this huge blob of data comes in, and the server says, help, we're under attack. It's like, no, this is just a big file that I'm sending.

Leo: Now, that's a good speed test. Too many packets.

Steve: So which one do you like?

Leo: I use the one at DSL Reports. You know them, of course, they've been around for ages.

Steve: Ah, yes, of course.

Leo: They have a very elegant one that doesn't use Java or Flash, which this one does, Speedtest does. And it's done in JavaScript, and it's really, I think, very nice. Gives you buffer bloat information. I don't know how accurate that part is, but I find it to be a - I think you'll like - when you see it, it's a geek's version of Speedtest. I think you'll like it.

Steve: Nice.

Leo: Yeah. DSLReports.com. It's the second - you know, just click the speed test.

Steve: And they've got a bunch of good utilities.

Leo: They're good. But they're not immune to this thing that ISPs do.

Steve: Same problem. The only way to really know is to transfer a big file from somewhere where you know the sending end is not bandwidth constrained.

Leo: Right.

Steve: You know, for example, if there was something that was being served by a CDN, you know, a big hefty content delivery network, you can presume that it's going to be able to deliver it. The problem is many of those are deliberately throttled because, for example, they offer video. And they don't want to just give you this massive video file all at once because you may only watch the first few minutes and then go, eh, okay, this is not, you know, and you switch away. So it behooves them to deliberately throttle that so that you've only staying ahead of where you need to be, rather than giving it all to you.

So the bottom line is it's not an easy problem to solve. I had an idea years ago for a GRC utility that would do this. The problem is it wouldn't work well in a shared sort of a packet-limited, cable modem-style mode. But it was sort of clever. It only - it used bouncing off of the nearest router, and didn't actually require you to transfer data from some third-party location. It was able essentially to check the speed of the so-called last-mile link, you know, the final link between your ISP and you. So I may, you know, once SQRL is behind me...

Leo: Uh-oh. I hear that phrase "I may." Not good.

Steve: No, so I'm not being distracted.

Leo: "I may." No, no, no.

Steve: Okay. So iOS and iTunes and the Apple App Store got hit. What was interesting was the way this was discovered that I haven't heard anyone talk about. And that is that there was a sort of a modest developer whose handle is "realpg." And they were just experimenting with an app, a very simple iOS app that had no Internet functionality, didn't use the iCloud APIs, and was just sort of like a little starter app for iOS, using, of course, the Xcode Toolkit, which is sort of the whole toolchain that Apple makes available for developing apps.

And as you guys mentioned on MacBreak Weekly, or as Rene was able to fill in the details, the problem in China is that obtaining those libraries from Apple, for whatever reason, is really slow. And I have seen that said many times. So many Chinese developers get them from a local server. And in fact in this case it was Baidu that is, you know, has a good reputation, they're a major Chinese Internet company, and so they were offering the Xcode toolchain.

Now, Apple has measures to detect tampering. And the developer would have had to have those, all the developers would have had to turn that off in order not to be notified that there was a problem. But I'm not deep enough into this to know whether there's, like, if there would have been other reasons to turn that off, like maybe there are other restrictions that that creates. But what this guy, this "realpg" Chinese developer found was that, although their code had no iCloud APIs and was not doing anything on the Internet, their code was popping up a dialogue asking for the iCloud password.

And so they said, wait. We don't want that. Who's asking for that? And then they used an iPhone that they use for development which allowed them to verify connectivity, and they found that there was a conversation with a remote command-and-control server happening when their program was being run. So that immediately told them that something fishy was here.

Now we know that I saw a number as high as 300 apps, Chinese apps, got into the App Store. Palo Alto Networks here in the states identified 39 different infected apps, and among them was WeChat, which is super popular. So you do have to wonder how a high-profile, high-end app like WeChat could have been compiled with sort of second-hand tools. But I guess that's just, you know, no one was presuming that the Apple toolchain could have a problem, although it also required disabling this notification that would have let people know that there was something fraudulent about the toolchain.

So anyway, all of the apps that were affected have been removed. And of course the problem now is that the word may not have gotten, you know, these were - some of them were apps specifically planted to be downloaded and to be sort of like phishing apps. WeChat would of course immediately fix their problem and push out an update and update all their users. There are many of these others that don't want to update anybody. They want the people to continue using this bogus app because it does have remote access capabilities. It's able to phish for user credentials. It's able to hook URL schemes. The scheme is the thing like the HTTPS.

It turns out that the iPhone is full of these schemes. iOS uses these schemes for all kind of interapp hooks. And so this thing would be reregistering schemes for itself, essentially taking them away from apps that used to have them, as a means of inserting itself into the control chain fraudulently. And it's also able to read and write data to the user's clipboard, which can be very dangerous because the clipboard is a protected resource in iOS where users have to deliberately copy and paste to and from it. This allows the app to have access to it because it's got special privileges, thanks to being part of the developer toolchain. And of course people often use the clipboard to copy a password from a password manager and then paste it into the web app or some other app that they're using. So the clipboard can often have sensitive information.

So anyway, this was a big breach in the, I mean, I wouldn't call this Apple's fault except maybe they need - I wouldn't be surprised if they react to this by somehow making it more difficult to be developing apps with a fraudulent toolchain. At least they were able to find them and quickly expunge them from the store.

Leo: That would be easy to do. Just put a hash out; right?

Steve: You would think so. And it may be as simple as what they're doing. I just, again, I don't know the ins and outs of app development, so I don't want to guess. But still, I wouldn't be surprised if Apple, you know, Apple generally responds to this kind of thing. And this was a big black eye, so they want to get that fixed.

I want to also mention, while we're talking about iOS, that there were so many security improvements in iOS9 that I can't even summarize them. In fact, I started counting them. And I thought, okay, well, this is dumb because, I mean, it just - across the board, WebKit had CVSes, just enumerable CVSes. So we should not, you know, the Security Now! message here is don't think of this only as offering some features that you may or may not feel like you want. For example, there's a new OpenSSL carried by iOS9 which

you definitely want because OpenSSL has been subject to a bunch of problems.

And so certainly Apple sends out security updates from time to time. This iOS9 had a boatload of, like, across the board, I mean, a breathtaking number of them. For anyone who's interested, I do have the link to the enumeration of them. If you want to kind of chuckle, Leo, you can bring up, it's on that second page here, the iOS9 support.apple.com and then a URL. But, I mean, it just goes, as you scroll, it just goes on and on and on and on. And many of them have, like, multiple CVS numbers. The WebKit ones especially are just like, you know, line after line after line of these. So lots got fixed in iOS9, definitely more than just, oh, here's some new features to support new stuff.

Leo: You know, this is kind of interesting. They give credit if there's a bug report to the people who discovered it.

Steve: Yeah.

Leo: Like here's a guy from the Yahoo Pentest Team, and somebody from Safeye, and Grayhash. That's kind of interesting. I guess that's probably a badge of honor, if you get your...

Steve: Yeah, and look where your scroll thumb is.

Leo: I'm only on "I." It goes on for quite a while. Wow. But this is a major upgrade. And in fact this is specifically a bug fix upgrade, so...

Steve: Yes. And of course we've got all of the new features in iOS9, as well.

Leo: Right. It does show you that Apple kind of stores up fixes.

Steve: Well, and I think probably there are some which, just because they are so deep in the system, there were a bunch of core networking fixes, I noticed. So there may be some that just don't lend themselves to incremental changes, where they had to, like, make some core plumbing fixes. We've talked about that in the past, where they would do a quick fix to solve the problem, but the real fix is going to take a lot longer. And so, you know, it just, like, involved rewriting some fundamental function which they couldn't break in an update, but they had to quickly patch it so that it wouldn't be exploitable, yet they really had to take their time and think something through more deeply. So this really has that sort of feeling to it, like this was a major set of, like, next step forward.

We did get, speaking of major, although I wouldn't call it as a step forward, anything to deal with Adobe Flash is a step backwards, but there was a major critical update. The good news is, of course, Chrome and IE are both managing Flash now on their own. Firefox still sort of waits for the user to give it a nudge. So for Firefox users you need to go to Tools, Add-ons, and then the Plugins tab under Add-ons. And there's sort of a little obscure link at the top. I wish they'd made this a bigger button, but it's just a little text link, and you wouldn't know it was a link except that it's blue, that says "Check to see if

your plugins are up to date." So you click that, and wham, you'll get a big, red, finally, whoa, Flash is bad, get a new one. And then you have to, of course, restart Firefox.

Maybe the problem is that I don't restart Firefox on a daily basis. I just - it's the first thing I run when my OS starts and the last thing I shut down when I'm shutting down.

Leo: Same thing with Chrome. You know, I may not restart Chrome for days.

Steve: Yeah. Although I do know that Google is updating all the time. I see like five processes running in my process list. So even if I'm not using Chrome, there are little gremlins, at least in Windows, that are running around keeping things updated for the next time that I do run Chrome. So you want to be at v19.0.0.185. I was at 18.something. So now we're at 19.0.0.185. And in case of Firefox, you may need to give yourself a nudge. And they fixed nearly two dozen vulnerabilities.

And this is important because, as we'll see, one of the mentions, when we get into talking about content blocking, just today was a blog entry about the Forbes.com site serving malware for about a week. And it was exploiting a collection of seven different Adobe vulnerabilities, Flash vulnerabilities in one ad. Just like, if there was anything that you had missed, this thing would burrow into your system just by visiting Forbes.com. So while...

Leo: Terrible. For a week?

Steve: Yes.

Leo: I mean, everybody goes there.

Steve: Yeah. So it was - I did also want to mention to remind people that, while you're in Firefox, if you're a Firefox user, when you're at the plugins page, or I'm sorry, I think it's at the, yeah, at the Add-ons Plugins page, on the far right is the Ask to Activate. You definitely want to do that. So, for example, Leo, you were mentioning that Speedtest.net is a Flash app. And sure enough, you know, when I go there on Firefox, I get just sort of a gray rectangle, and it shows a weird little Lego block, and it says "Click here to activate." And when I do, then I get a, over by the URL, I get a pop-up from Firefox that confirms I want to allow, and then I can choose to allow and remember if I want to whitelist this site's use of Flash.

Anyway, so lots of protection. And we've talked about this in Chrome. Chrome is defaulting now, not for first-party Flash content. So, for example, if you were on YouTube, and for whatever reason you couldn't use the HTML5 player, and it did a fallback, then it would still play because that would be coming from presumably first-party content. But ads coming from a third party, those are now click to run. So probably something is show, if it doesn't run, but at least you're not running the active content, which is what enables the exploit.

So, you know, we're finally beginning to see, by default, these constant, year after year after year, vulnerabilities being fixed proactively by the way that they're being used in the system by the technology that non-listeners to the podcast, most people just use

their browser and go to sites, and they're not, you know, as I've often called it, the tyranny of the default. They're not tuning and tweaking anything. They're just hoping to surf the 'Net and not get hurt.

Okay. Now this was chilling. There was a lot of controversy generated by this. AVG, the number three most popular antivirus tool that offers a free version which many people use. Number one most popular is Microsoft. Number two is Avast. And number three is AVG. They've changed their privacy policy, notified their users that a new policy becomes effective as of October 15th, so about three weeks from now, really raising hackles within the privacy community because the new policy states that users of the AVG antivirus will be permitting AVG to sell search and browser history data to third-party advertisers in order to make money from its antivirus software.

Now, the problem is that AV software inherently runs in our system with elevated privileges so that it's able to detect and block malware, adware, spyware, and other threats. And AVG is one of the AV suites which we've discussed in the past that installs its own root certificate into users' machines, like Superfish was caught doing, specifically to allow it to intercept, decrypt, and inspect all web browser traffic. So we're not saying anything nefarious is going on, but essentially what they've decided is that they, too, are sort of going down this path of monetizing the habits of their users.

Now, what's really confusing is that an AVG spokesman explained that, quote, "any nonpersonal data collected and sold to advertisers would be cleaned and anonymized, making it impossible to link it back to individual users." And then the spokesman said, "Many companies do this type of collection every day and do not tell their users." So they're trying to claim some cred for, like, being right upfront with the fact that they're going to start monetizing by selling their users' search and browser history data.

Well, the problem is it's only valuable to advertisers if there's some way to tie it to users. So while AVG may be sanitizing it in some way, it must be that things like the cookies of the queries, which may not contain overtly user identity data, is the token that the advertiser has previously assigned to this person in order to associate them. Meaning that selling data that couldn't be tied back to the user wouldn't generate any revenue.

So it's got to be that essentially, when you're using AVG in the future, essentially they're now tracking you and saying that they're going to monetize this and sell this to advertisers. Which just creates another pipeline for this kind of information, you know, browsing and search history is what they're talking about, to sort of this unseen advertising tracking database facility. So I just wanted to make sure that our users who may be using AVG were aware that this was something that was going to be happening starting the beginning of next month.

Okay. Cisco. Guys at FireEye discovered in a handful of routers, I think like 14 or 15, located in four different countries, and these are, when I say Cisco, people can sort of now think of Linksys because, you know, they bought Linksys. But no. These are like big iron Cisco green. Green is the color of Cisco routers, always has been. That's the color that they chose, sort of a nice forest green. And you can walk into a datacenter and immediately spot Cisco network equipment because it's Cisco green color. So, and there were three different model numbers of affected routers where they found this stuff, but those are because they're the big iron style routers.

So these are typically positioned on major networks' borders. They're the so-called "border routers" which are used as the first line of contact between the Internet and the internal Intranet that links the network to the Internet. What they found was, they called it "SYNful Knock" malware, S-Y-N-F-U-L, and they called it "knocking" because the way

this works is very clever. We've talked about port knocking of various sorts in the past. The idea is, with port knocking, you can have a router, like we all have NAT routers now on our borders. You can have a router which silently drops any unsolicited incoming traffic. That's a beautiful firewall-ish nature of any NAT router is that, when someone is trying to ping you or probe you or scan you or whatever, and this thing comes from a source IP and port that the router didn't first see traffic leaving your network bound for, so that the return traffic is expected, if it's just coming from something that the router's never seen before, it ignores it.

Well, it's possible for the router to log that, or to make some note of it. And then the idea is that some remote IP, which would normally be blocked exactly in that fashion, could deliberately send packets in a specific sequence to a specific set of ports. That's called a "knock." That is, you know, like knocking on the door, if you don't get in, but you knock. And so if the router is monitoring that and sees a specific set of packets arriving in a certain order on specific ports, that's sort of - it's a weak form of access, "weak" because of course an outside observer could see the knocking, learn the knock sequence, and then themselves gain entry. So it's not super secure. But it's an interesting kind of cool technology, sort of a side effect of what the border router is doing.

These guys found that the Cisco firmware itself, Cisco has their own operating system called IOS, Internet Operating System, which had been modified to look at the TCP sequence numbers. Remember that the SYN packet contains a 32-bit sequence number which numbers the data that the TCP connection will subsequently be sending. So the SYN packets contain essentially pieces of information which are normally unpredictable and hopefully random. There have been lots, like a history of attacks when that was not the case, and we've talked about that in years past. Pretty much all of the TCP/IP stacks have that fixed now. But this can be used for knocking also.

So the firmware would be looking for specific sequence number of incoming SYN packets. And when they qualified - oh, SYN and ACK. Normally, I'm trying to think, I think the acknowledgment field on a SYN packet is just null on a SYN. Then on a SYN/ACK that is the TCP packet sent in response, the sender's own sequence number is in the 32-bit SYN field, and the acknowledgement field contains, I think it's the SYN that was received. That would make sense because then that's going to be advancing as essentially the acknowledgment field is the most recent, the sequence number of the most recent packet received from the sender. So the acknowledgment field says, okay, I've seen, with no gaps missing, nothing out of order up to this numbered byte starting with the byte that we began.

So these SYN packets came in with both of these fields, so essentially 64 bits of specific information. If they matched a pattern of some sort, then that would identify the IP of the sender as authenticated and essentially give them access to a command-and-control that was built into this firmware. So this is really bad. It was found in the Ukraine, in the Philippines, in Mexico, and India. So, and FireEye did not tell us which companies or networks these were associated with. They kept that quiet. Cisco's been notified. They've put out, made a fix available for this. And again, this is not, as far as we know, it didn't leverage a Cisco bug, but it did get into the routers. And the term, the NSA term that we learned when we talked about some of the early Snowden revelations, were "implants."

So the presumption is someone had brief physical access. There was some mention made in the FireEye disclosure that the routers used default passwords. So it may have been that someone accessed this as it was being shipped to its destination - we believe this is some of what the NSA is able to do - implanted these changes in the firmware. Then the router was shipped on to its destination. Then the router was configured and set up, but the firmware was just left alone. All the people had to do was re-update the - there's like

a ROM-based loader, sort of the pre-IOS loader. ROMmon, it's called, the ROM Monitor, that's where this stuff was located.

So just reflashing the IOS, which I've done on my Cisco routers through the years, wouldn't have removed this. This is down at a deeper level. So it required sort of high-level access to this. And it was then deeply implanted and very sticky. So, you know, this is what's known to be going on, and FireEye found some instances of this out in the wild.

Ashley Madison. It was believed that the use of very good Bcrypt - Bcrypt is one of the so-called PBKDFs, password-based key derivation functions. The idea being - the use of these is the state of the art for protecting passwords on a server, the idea being it's no longer enough just to hash the data, even with a per-user salt, and then use a good SHA-256 salt. The problem is using strong GPU, or now purpose-specific ASIC hardware application-specific integrated circuits, hashing is so fast now, really driven by the bitcoin mining phenomenon, that what you need is to do iterative hashing, so not just hash, but hash many times. Bcrypt is an algorithm we've discussed which does this. And in fact Ashley Madison was using a strength that required 4,096 iterations within the Bcrypt algorithm.

So the presumption was, of the 36 million passwords in the records that were obtained in the breach, it would take a long time, basically be infeasible to crack them. And there was some early success, only because the passwords were so bad, you know, 1234, 123456. And given the nature of the Ashley Madison site, one's imagination could lead you to some creative guessing, which turned out to be accurate in many instances. So the hackers were having fun guessing what people might have used as their Ashley Madison passwords. And what do you know, that happened to succeed for many accounts.

Completely separate from that, what was discovered among the data was that, of the 36 million accounts, 15 and a quarter million had some additional data associated with it. It had an MD5 hash. And no one's really sure why. My theory is that this was old technology; that the very, very first instance of the web software that ran Ashley Madison in the beginning used a simple MD5 hash. I mean, even then, they were doing some hashing, which is better than many high-profile sites that didn't even bother with that. But it wasn't very good.

And in fact it looks - the reason I believe this is that they were lowercasing the username and lowercasing the password. They were then concatenating them with a pair of colons and taking the MD5, the standard MD5 hash of that. Well, that would be a reasonable hash to use for verifying login like for your beta of the site, as you're getting going. And I say that because then they fixed it a little bit later, not very well. They added the email address and a supposedly secret bunch of gibberish, sort of like unchanging salt. You know, it was the digit 7 and then 3, an @ sign - shoot, I'm blanking on the name of that thing. Not a tilde, the little hat symbol.

Leo: It's got a hat.

Steve: No, it's the Shift 6.

Leo: Oh, a circumflex.

Steve: Circumflex, yes. I actually had the word, but it somehow didn't feel like it was a circumflex.

Leo: I love it. Okay.

Steve: The raise to the power of symbol in many languages.

Leo: Yes.

Steve: And then lowercase bhhs&#@&^@8@*S". Or I guess the double quotes is the end of my quoted string. So it's like, okay. So that's some salt that they added. But so the point is that the username is part of all of these MD5, is part of the hashing. And the username is available in the breached data. So that removes it. Then all we're left with in the first case is lowercasing the password. So of course that discards any case sensitivity from guessing, so the guessers only need to use lowercase guesses. The person's email address is also part of the observed data. The bottom line is that, as a consequence of this probably obsolete, that is, only, what, less than half of the accounts had this. My guess is this was used early on, and then it was replaced by the really good 4,096 iteration Bcrypt PBKDF function. But they left the original data there because of course you'd have to have some crossover. You need to allow your users to log in using their original username and password. So you'd still have to have the weak username and password authentication in place while you then generated the new stronger one. But if we know anything, we know that these guys weren't doing a very good job. For example, people were able to pay to be removed from the database. And we now know from the retrieval that they were never removed.

Leo: From anything.

Steve: Yeah. So neither were what looks like...

Leo: They were just scammers.

Steve: Yes, yes. Neither removed what looks like the early versions of their username and password authentication. As a consequence of the fact that that stayed around, it turned out to be vastly simpler to guess people's passwords. For 15.26, that is the 15 and a quarter million accounts that also had these weak MD5 hashes, 11 million have now been cracked.

And of course the problem is we now have username, we've got email address, and their password. So anyone who reused a password from some other site is in more trouble because we know their name; we have their email address; and, if they were among the early adopters of Ashley Madison, it's probable in at least 11 million cases their password has now been cracked because the MD5 token makes it easy to perform a brute-force password guess, especially when you're throwing away case. That then allows them to - then they verify that against the Bcrypt-verified password to essentially crack the user's data completely. And if you did reuse your passwords, then you're in trouble. So, yikes.

Leo: Yikes.

Steve: Yeah. The breach that just keeps on giving. Okay. So this one, this is not about security, but we've talked about the technology that is entering the automotive sphere. And apparently this affects just shy of half a million VWs and Audis. It's the Jetta model years 2009 to now, the Beetles from year 2009 on, the Audi A3 from 2009, the Golf from 2009 on, and the Passat, only 2014 and from now on.

This was from the EPA.gov site. They said the "EPA is issuing a Notice of Violation (NOV) of the Clean Air Act to Volkswagen AG, Audi AG, and Volkswagen Group of America, Inc., collectively referred to as Volkswagen. The NOV," this Notice of Violation, "alleges that four-cylinder Volkswagen and Audi diesel cars from model years 2009-2015 include software that circumvents EPA emissions standards for certain air pollutants. California is separately issuing an In-Use Compliance letter to Volkswagen, and EPA and the California Air Resources Board have both initiated investigations based on Volkswagen's alleged actions."

As described in this Notice of Violation, "a sophisticated software algorithm on certain Volkswagen vehicles detects when the car is undergoing official emissions testing, and turns full emissions controls on only during the test. The effectiveness of these vehicles' pollution emissions control devices," they write, "is greatly reduced during all normal driving situations. This results in cars that meet emissions standards in the laboratory or testing station, but during normal operations emit nitrogen oxides, or NOx, at up to 40 times the standard. The software produced by Volkswagen is a 'defeat device,' as defined by the Clean Air Act."

And then there was some interesting commentary that I appreciated, that provided a little bit more details. This writer posted: "Just when the promise of diesel was gaining new legs thanks to lower diesel fuel prices, new low carbon fuels, and engines that could easily meet the strictest emissions standards in the world, the EPA issued a Notice of Violation letter to the largest seller of diesel engines in the U.S. and the largest auto manufacturer in the world.

"During dynamometer testing for both Emissions and Fuel Economy certification, 2.0L VW TDIs" - and I didn't have a chance to look up that acronym. Do you know what that means?

Leo: Oh, I should because I have an Audi. But I don't have a TDI engine.

Steve: Oh, so, okay, so it's a type of...

Leo: It's their diesel, yeah, it's, you know, it's probably a marketing term, but I'll look it up real quick, yeah.

Steve: Yeah. "...VW TDIs from 2009-2015 worked in a..."

Leo: Turbo-charged Direct Injection.

Steve: Ah, nice. So nice high-end diesel.

Leo: Turbo diesel, yeah.

Steve: Yeah.

Leo: Although these were only four-cylinders that they detected it in, so...

Steve: Right, two-liter. So that's the smaller of the line; right?

Leo: Yup.

Steve: Yeah, "...worked in a dyno-mode tune that performed satisfactorily and easily met the legal limits for a variety of pollutants including NOx." Okay. So this guy said: "Unfortunately, when the ECU detected that the car was not on a dynamometer" - and this writer says "I am still not entirely clear as to how this was accomplished," and that would be a juicy little tidbit - "the emissions controls were allowed to relax, providing both the power and efficiency" turbo diesel injected, oh, wait, turbo, yeah, turbo diesel injected, "TDI owners have come to rely on. This in turn allowed between 10 and 40 times the legal limit of NOx to be emitted while driven in the manner that we all drive.

"In an EPA teleconference call with the media, Cynthia Giles, an enforcement officer at the EPA, stated: 'Put simply, these cars contained software that turns off emissions controls when driving normally and turns them on when the car is undergoing an emissions test.'"

And then this guy had a little more information, saying: "The defeat device, as it was called, or dyno-mode programming within the Bosch EDC17 ECU, has been a part of OEM testing for decades. It was the relaxing of the emissions control systems code that caused the real problems."

Leo: Eric Duckman suggests that the way you would tell if it's on a dynamometer is if it's going road speed RPMs, but nobody's touching the steering wheel, there's no movement of the steering wheel.

Steve: Ahh. Interesting.

Leo: Yeah, that would be one way, for sure. Or no weight in the seat.

Steve: Right, because they probably...

Leo: All of these cars have weight in the seat sensors for seatbelt.

Steve: Right.

Leo: There's probably a number of ways to do it. I mean, it's such an abnormal, you know, you get all these high RPMs, but nobody's in the vehicle. That would probably be a giveaway.

Steve: Yeah, that might. And I wouldn't be - also don't they, I guess it's - no, no, I was going to say they're also not level normally. But I think they are level. They merely put the back wheels on those two rollers, yeah.

Leo: Yeah. You know, there's a DMCA angle on this because the DMCA forbids reverse-engineering software in these vehicles. Had that been - and, frankly, a lot of researchers would have liked to have gotten in there.

Steve: Yeah.

Leo: Had that been legal, we might have seen this a long, you know, we might have noticed this years ago. But it's not legal to reverse engineer.

Steve: So what do you think this means? Do you think this means that the cars cannot deliver the power that they're selling and also meet the strict California emissions standards?

Leo: Exactly.

Steve: Wow.

Leo: That's exactly right. What people don't like about diesel is the hesitation. Diesel has a lot of torque at startup, but they don't like the hesitation. And I bet you anything it has to do with just trying to get the, you know, that was what they were selling these diesels with. Oh, these new diesel engines, they perform like gasoline engines.

Steve: They're so - they're small.

Leo: They're efficient and small. Not a lot of incentive to buy them here. Diesel fuel is more expensive than gasoline in California. But, yeah. I wasn't tempted. I have an Audi. It's shameful. I won't buy another Audi.

Steve: Get caught. Okay. Get this now. Oh, boy. This is the short version, is how to do it all wrong. This is India's, I'm not kidding you, there's a PDF link in the show notes to the whole draft, I've excerpted a few tasty bits, but mostly I've summarized this. This is India's draft national encryption policy. So we've talked about how we're living in

interesting times. You know, I, years ago, famously suspended work on CryptoLink because I was worried about what was going to happen. And we've got now Apple fighting with the DOJ because the DOJ wants iMessages to be turned over. Microsoft fighting because the DOJ or law enforcement wants email that's being stored in servers in Ireland and so forth.

So India has produced a request, a for-comments draft policy. PublicIntelligence.net summarized it, saying: "The following draft copy of the National Encryption Policy was released for public comment by the India Department of Electronics and Information Technology. The policy has been widely criticized for requiring businesses, Internet service providers, and even private citizens to store decrypted versions of encrypted communications for 90 days to provide to the government and law enforcement," and I'm adding, you know, if required. So that's sort of the - that's the crux of this is that their solution to the problem, such as it is, or that is the solution, is to say, okay, one way or another, encryption must be made available in plaintext form for 90 days from the time it was created, transmitted, stored, whatever.

So first of all, this applies to both data in motion, in transit, or at rest. And, you know, it starts off with your standard, you know: "The recognition of the need to protect privacy and increase the security of the Internet and associated information systems have resulted in the development of policies that favor the spread of encryption worldwide." Yeah, right. "The Information Technology Act 2000" - so apparently there's some existing boilerplate that they've got from 15 years ago. "The Information Technology Act 2000 provides for prescribing modes or methods of encryption (Section 84A) and decryption (Section 69). Taking into account the need to protect information assets, international trends, and concerns of national security, the cryptographic policy for domestic use supports the broad use of cryptography in ways that facilitates individual and businesses privacy, international economic competitiveness in all sectors including government." Huh.

"This policy is not applicable to sensitive departments/agencies of the government." So they're making an exception for themselves and national security. So it's like, this is what we're going to impose on businesses and individuals, but not on the government, which performs "sensitive and strategic roles. This policy is applicable to all central and state government departments, including sensitive departments and agencies while performing non-strategic and non-operational roles, all statutory organizations, executive bodies, business and commercial establishments, including public sector undertakings and academic institutions and all citizens, including personnel of government/business performing non-official/personal functions."

So just to summarize, they said, under Objectives, the first and main objective: "To synchronize with the emerging global digital economy and network society and use of encryption for ensuring security and confidentiality of data and to protect privacy in information and communication infrastructure without unduly affecting public safety and national security."

So basically, and I'll just summarize, What they are imposing is full regulation of all use of encryption within India's borders. The government will dictate under this the allowable encryption algorithms and key lengths, it says, as I mentioned, both data in transit, so communications links. They did accept the web use of HTTPS. So that they're not going to - that they're not targeting. It did say TLS and SSL for Internet. So they're not expecting to be able to break that. But, for example, iMessage, or any encrypted chat, any encrypted conferencing like we're using right now. I mean, basically everything. And they're saying that you could only use approved-of algorithms and key lengths.

And mostly the way they're imposing this is to say that, upon demand, plaintext must be producible for anything encrypted within the past 90 days. The actual verbiage is: "On demand, the user shall be able to reproduce the same plaintext and encrypted text pairs using the software/hardware used to produce the encrypted text from the given plaintext. Such plaintext information shall be stored by the user/organization/agency for 90 days from the date of transaction and made available to law enforcement agencies as and when demanded in line with the provisions of the laws of the country."

Anyway, so I won't go any further. I had a bunch more stuff in the notes, if anyone is interested. Basically, it's an attempt to fully regulate the encryption problem, to solve the problem by regulating it within India's borders. Oh, "All vendors of encryption products or service providers offering encryption services shall necessarily register their products and services with government for conducting business in the country."

Leo: So they just withdrew it.

Steve: What? Oh, no kidding.

Leo: About an hour ago.

Steve: Okay.

Leo: But the real issue, it's so funny, it seems to be that the complaint was about WhatsApp, was about messaging.

Steve: Wow, a high-use messaging app.

Leo: Right. The government says the rules weren't meant to target WhatsApp, shopping, and other common activities. This is Ravi Shankar, the fabulous sitar player and Telecom Minister. But I imagine they'll be back with another one. This comes from The Times of India.

Steve: Well, so what we saw was what their intent was. But if they're intending to exclude something like WhatsApp, then that would sort of presume - I guess, then, what is it that they want to get? Obviously they want to be able to get whatever they need. But WhatsApp and iMessage would seem to be sort of then within the same domain.

Leo: Pavel Durov, the guy who created Telegram, which as you know has this sort of security model...

Steve: Yeah, sure, sure.

Leo: Was asked at TechCrunch Disrupt yesterday, well, it seems to be the case that

ISIS is using Telegram.

Steve: Oh, no. Is that really true?

Leo: And he didn't deny it. He said, well, you know, the problem is that people deserve privacy. And you can't - and Phil Zimmermann always said this about PGP. You can't - ISIS uses telephones. You can't eliminate a technology just because a terrorist might use it. And there's an overwhelming need for private communications for normal people, as well. So it was an interesting conversation. By the way, currently 60 million active users of Telegram. They just crossed, last year, in fact early this year, there were about a billion messages a day. He said they're doing 12 billion messages a day. Not a lot more users, just a lot more use.

Steve: And half of those, Leo, are because you love it. That's right.

Leo: And I love it because of the stickers.

Steve: I know. You've got them all.

Leo: It has nothing to do with security. I like the stickers. So that was a quick victory over the Indian government. Actually, I don't think that's over. I think we'll see more. But it's interesting. They were obviously stung.

Steve: Well, yeah. It's why I brought it up. We knew, yeah, we knew that it was, like, you know, this is the first formal proposal that we've seen. I don't know how it's going to sort itself out. But I'm glad that this got brought up at Disrupt, and that the Telegram guy said, yeah, you know, we provide encryption because our users want it. Everyone has a right to not be listened in on. And we're sorry if bad guys use it. You know, you'll need to get a hold of their conversation by taking over their phone, rather than by...

Leo: Yeah. There are means. There are ways to do it; right?

Steve: Yeah. Yeah.

Leo: He said privacy is - here's the actual quote. This is Pavel Durov: "Privacy is ultimately more important than our fear of bad things happening like terrorism. If you look at ISIS, yes, there's a war going on in the Middle East. Ultimately ISIS will find a way to communicate with its cells. And if any means doesn't feel secure to them, they'll just find something else. We shouldn't feel guilty about it. We're still doing the right thing, protecting our users' privacy." Seems right.

Steve: Yeah. Yeah. I think that's the way it's going to have to shake out. Unfortunately, as we've discussed, no one can see a way of giving law enforcement what we would be

willing to give them if we could without sacrificing the integrity of the technology. But there just doesn't seem to be a way. Yes, we could build an extra set of keys into everything. But then it would just, I mean, then we'd have nothing because look at how poorly the federal government is a caretaker of the secrets that we've given them. They're just not. So.

So. I did want to mention that the StageFright fixes continue to drift out. I get tweets from people whose random Android phone receives an update, and they go and check with the StageFright detectors, and, yay, they get all green. In some cases, they're getting that - there's only that one remaining one that came out later. That's still vulnerable. And it's just because there's a delay through the pipeline getting the updates made and fixed and merged into the firmware and then pushed out. But even as late as a day or two ago I saw somebody just updated their phone or had their phone update and was now StageFright safe.

Leo: Yeah, all my Android phones are now safe, according to the Zimperium test, which is the one that seems to find that last little bug.

Steve: That's the one you want.

Leo: Even the Galaxy Note 5, a Samsung device, that was updated. You know what's not updated yet, but it's brand new, maybe it will be soon? The Moto X Pure, which I just got last week, is still not updated. All the Nexuses are, though. They've been very - OnePlus 2 has been for some time. So that, yeah, it's getting out there.

Steve: And I would say that this is, given that it's known now, and we know that blocking the auto fetch of MMS messages creates a barrier to easy exploitation, this gives us a benchmark for the way the various providers are responding because there isn't a provider that should not respond. And everyone can see now how long it has taken until they've had the StageFright updates. So it's a little bit of a canary at this point, where we can say, hey, look, you know, as you said, Leo, all of yours, with one exception, are now updated, and that one seems to be pretty new. So they're just probably scrambling around, pushing the rest of them out the door.

I did get a tweet from a listener. I just wanted to share it. Simon is his name, @sphere_au, so maybe he's Australian. Anyway, he said: "Just got @letsencrypt working on my web server." And he had a shout-out thanking someone for the support in the forums for ironing out some little config issue. And then he said, "ping @SGgrc." Speaking of which, I'm closing in on 50,000 followers in Twitter. So I'm at 499 something.

Leo: That's nice.

Steve: So it'll be fun to be north of 50.

Leo: @SGgrc, yeah.

Steve: @SGgrc. Couple miscellaneous things. I wanted to say, Leo, I listened to you and Megan yesterday on iOS Today. Toward the end of the show, the topic of batteries came up. And you were fabulous.

Leo: Oh, was I right? I hope I was.

Steve: Absolutely.

Leo: I quoted you.

Steve: I was holding my breath, and everything exactly right.

Leo: But why does Battery University say other things about it?

Steve: The only thing I can think is that there is one danger, and that is that the state-of-the-art lithium-ion cells really get unhappy if you overcharge them.

Leo: Yeah, but no modern - no decent device allows that.

Steve: That's exactly correct. So they had some advice that you were exactly correct to challenge and criticize them on, which was, you know, don't leave it in, don't leave it plugged in overnight. It's like, what? And don't charge it to 100%. What? Because, no. Everyone is going to do both of those things. And as you said, it used to be that - you know the RC model car chargers, where they charge these batteries super fast, they get really hot, but they just want to recharge them immediately, stick them in the car, and run them around again. Those are NiCads.

And NiCads have an interesting phenomenon, and that is, you can charge them super fast. Their voltage peaks and then starts to drop. And so what the fast chargers do is they look for that, they look for a drop, the beginning of a drop in the battery terminal voltage. And the second they see that, they stop. And that allows you to just cram the current back into the battery, bring it right up to a full charge in just a few minutes, and then stick it back in your car, and off you go.

The whole deal with lithium-ion batteries is, as we know, very different. We wish they were fast to charge, but they're not. You need to charge them slowly. And what you do is you very carefully monitor the terminal voltage and just stop when it reaches like a known full charge. You have to be very careful with them not to overcharge. All those fires that we've heard about, like the laptop explodes in the user's lap, well, it was plugged in and charging, and the charging circuitry didn't cut off. And, boy, if you overcharge these, they're not happy. But as you said, Leo, we've solved this problem. It's like, this has been solved. So but you were, again, your advice was...

Leo: Well, let's repeat it because people are now wondering what's the advice.

Steve: The advice is lithium-ion really dislikes deep discharge. I had a buddy who complained. He lives in the Philippines and visits twice a year. I met him at Starbucks, and we've become friends. He was saying that he was annoyed because his phone - so I happened to see that his phone screen had, like, red on the battery. And I said, "John, what is that about?" And he says, "Oh, yeah, it's almost dead. I've got to plug it in." I said, "Don't you keep it charged?" He says, "Well, I charge it whenever it almost dies." I said, "No, no, no, no." And then he tells me that he keeps, like, burning out his phones and his iPads. And I said, "Yes, you cannot deep discharge lithium-ion. They really want to be kept near charge." So charge it at work. Charge it at home. Do not deep cycle them. That's the main advice.

Leo: Which is funny because you hear that advice, deep cycle them. So the two rules of thumb are, well, you need to know that every lithium-ion battery has a certain number of charge cycles. That's fully charged, fully discharged. And the rule of thumb is keep it plugged in because you can have fractional discharge. It all adds up, though, and you don't want to have, you know, you don't want to hit the number of cycles. And then if you're storing it, I think you said keep it at half; right?

Steve: Yes. And you're completely correct again. It was perfect.

Leo: Well, I memorized what you said.

Steve: So the one reason that it might make sense to occasionally do a discharge is - and it's referred to as "battery conditioning" sometimes. What that actually is doing is synchronizing the device's awareness of how much charge is in the battery.

Leo: Right.

Steve: Because if you just sort of float the battery around, plugging it in, pulling it out, plugging it in, pulling it out, it's possible for - because the battery itself doesn't provide any indication of how much charge it has. What happens is, in a lithium-ion, the terminal voltage stays flat for a long time, drooping only very slowly. And then toward the end it falls off.

Leo: That's why it's at 100% sometimes for a long time, and then it starts to go down. That makes sense.

Steve: Right. And so what's happening is a software algorithm is having to guess. It's sort of running a down counter while it sees that you're using power, and it's not plugged in. It's kind of going, okay, well, I may be 89, uh, maybe 85. I mean, it has no idea. So it's just sort of - it's just ad hoc. When you fully charge it, then it can say, okay, we're back at a hundred. But it still can't quite tell where the bottom is unless you reach it. That is, because as you also properly said, there is a total, like, there's a total amount of use that lithium-ion batteries will have.

But so an older one will, you know, it'll only hold maybe three quarters or two thirds as much energy. But again, the counter doesn't know unless it watches it go down and run

dry. Then it goes, oh, now I'm recalibrated on the current state of health of this battery. So then you charge it fully, and now it knows. It's a little bit more pessimistic, but more realistic, about how much time is actually left.

If anyone, for example, has had a device suddenly die when it said, wait a minute, you said I had 50% left, and suddenly it's dead. It's like, yes, that's a perfect example of the device not having been able to calibrate itself. So sometimes it is useful to bring it all the way down, but not as your normal daily, I mean, some people don't charge them until they're dead. And it's like, oh, boy, you know. And unfortunately, we don't have replaceable batteries in these devices any longer, so it's like a problem if you kill it. So 100% advice.

I just did want to mention, I heard you talking, I guess it was about Audible, of course, the other day. And I have to tip my hat to Peter Hamilton, who's one of our favorite sci-fi authors. It was the middle of the night. I was doing some reading. I was changing books that I was going to read. I've been studying stress for the last year, actually, and anxiety, and the whole hormonal and brain mechanism systems. THE reference text for this is by a neuroscientist. The title of the book is "Why Zebras Don't Get Ulcers." It's a dumb title for a fabulous, I mean, deep science book. But I was reading "The Cortisol Connection," and I was switching back to "Zebras," and I thought, I'll just kind of take a look at "The Abyss Beyond Dreams." That's Peter's latest pair of books. He was going to do a trilogy, and he decided just to make it a - what do you call a two?

Leo: Duology?

Steve: Duology? Yeah, I guess.

Leo: I don't know. I guess that would work.

Steve: Anyway, I just - so I just started, like, I'll just kind of see what it's looking - and I just immediately got sucked in. Now, I didn't lose control. It was late, and I just - I stopped. But I just said, wow, you know, I would be reading it except that he is not finished. "The Night Without Stars" is the follow-up to "The Abyss Beyond Dreams." And he got me with "Pandora's Star," so that I was so annoyed when I finished "Pandora's Star" because it was only the first half of the story. And we were all like, oh, what happens now? In fact, I had to reread "Pandora's Star" when he finally came out with the second half of that story.

So I'm not doing it again, Peter. I'm waiting until "The Night Without Stars" is finished. He's off on some child's fantasy book series. And it's like, come on, come back to hard sci-fi. We need a second of these because this the Chronicle of the Fallers is the series that these two books are part of.

Leo: And I'm putting that Sapolsky book in my Audible cart as we speak.

Steve: Is it available on Audible?

Leo: Absolutely, "Why Zebras Don't Get Ulcers."

Steve: Ah. That's the one, baby. It is really good. I vouch for that from a - the science of stress. Really interesting.

Leo: Well, I've got a lot of it. So I'll go read this.

Steve: Listen to it. I think it'll help a lot.

Leo: Yeah, yeah, yeah.

Steve: I did get a nice note. This is sort of a different kind of testimonial from a Jim Gerry in League City, Texas. He says, formally, "Mr. Gibson, I have been a loyal customer for a LONG [all caps] time. How long? I was digging around my closet and found an old 'SpinRite: A Guide for Owners, Version 1.2.'"

Leo: Wow.

Steve: And we actually used to publish a hardback book that had a spiral, a metal spring binding, it was very expensive. But, you know, I wanted it to be classy. And he said, "And that's what prompted me to write this email. I just wanted to say thank you! I have used one version of your software on every hard drive I have ever owned, going all the way back to my 10MB hard disk drive on my IBM PC/XT. I cannot tell you how many hard drives I've owned over the years, but your software was used to rescue my data time and time again, after the drives inevitably quit booting. It was always easier to recover my hard drive with SpinRite than to restore from backups. Your software is as relevant today as it was in the 1980s. Simply amazing." So, wow, Jim. Thank you...

Leo: Isn't that sweet.

Steve: ...for the shout-out for SpinRite.

Leo: All right. I am all ears for this because I turned on, with my iPad Mini, I turned on content blocking, and it really does make a big difference. But I just bought the first one I saw, which was Crystal.

Steve: Right.

Leo: So I'm very interested in what the differences are and so forth.

Steve: Yeah, and in fact Crystal is the only one that I don't recommend.

Leo: Oh.

Steve: Although, well, only because there is no whitelisting feature.

Leo: Right.

Steve: And I think it's entirely...

Leo: It's too simple.

Steve: Right. And again, this is the first week.

Leo: Right.

Steve: So it would be trivial to add that. So I wouldn't - I don't want to blackball it, just because it doesn't have that. But I think it is entirely foreseeable that we're going to enter a phase where we as browsers, browser users, browsing sites, are confronted with a choice. The site says we're ad supported. You're blocking our ads. Therefore we're blocking you. Some sites may just say, hey, you know, we're supported, so please consider lowering your shields so that your presence here can support us because that's what makes this go. Others may just say, eh, you're not getting in. You know, basically the equivalent of the paywall, but in this case an ad wall. So what that says is that we need to be able to whitelist.

Now, remember, though, that any time there's a problem with looking at a site in iOS, you can hold down the little refresh curlicue, the little refresh arrow, and up pops a popup that allows you to reload the page without filtering. So it doesn't make it sticky and, like, whitelist the site for future visits. But it just says, oh, if you want to, you can view this page with any content filtering which you may have in place disabled, just for now. So that's a workaround for not being able to whitelist. My sense is, for example, sites that you want to support, like iMore, you're probably more likely to say, yeah, I want to put them on my whitelist so that I can support them. Anyway, so that's one of the criteria. Interestingly, Crystal was the only one among those that I looked at that did not offer that feature. Anyway, we'll come back to this.

You might want to look at Lifehacker.com, Leo. Halfway down the page is an article, "The iOS Adblockers That Speed Up Your Browsing Most." So you can bring that up in a minute. And I went to Lifehacker.com to the home page. About halfway down, if you pull the scroll bar down halfway, you'll find the link to that page. And there is a chart there that shows the result of their benchmarks.

In the meantime, our friend Dave Winer, who we talked about a couple weeks ago - yeah, you found it. And this is not Dave "Whiner," although Dave does tend to whine a little bit. Of course he's the father of computer-based outlining that I am a disciple of. I love outlining for organizing things. And as I understand it, Dave invented blogging; right? He was the father of...

Leo: He's widely, yeah, I don't know if invented blogging, but he created the RSS feed.

Steve: Okay, right.

Leo: He also is in some ways the father of podcasting because he was the one who extended RSS feeds to include binary attachments.

Steve: Ah, nice.

Leo: And without which you have no podcasts.

Steve: No podcasts.

Leo: So Dave is definitely a guy. He's the guy. I mean, I think his outliner - did he do ThinkTank? I know he did MORE a little later on.

Steve: Oh, I think ThinkTank for sure.

Leo: ThinkTank was the one.

Steve: Yup.

Leo: And MORE was his...

Steve: I used ThinkTank on an Apple II back then and was happy when it made it over to Windows.

Leo: And as I mentioned last time, he has an online outliner you can use for free.

Steve: Yes.

Leo: He's really - he's cranky, as we all are.

Steve: Yup.

Leo: He's seen a lot. But I love him. He's great.

Steve: So to put this in context, this was a blog, his blog on Scripting News, it's Scripting.com. So you get it. If he's got Scripting.com, he's been around for a while. And he posted this on September 19th, so that was Saturday. So iOS9 hit on Wednesday. So this is, what, three or four days into Adblockalypse that Dave writes what I'm going to share. And the reason I want to do this is this is what led me to a question that I hadn't thought to ask before.

So he wrote: "I think of advertising as 'unwanted commercial messages.'" He says: "The unwanted part is key. I do a lot of seeking of commercial information using the web. We all do, all the time. That's how business works on the web. It seems to me that news orgs have to figure out how to make people come to their sites seeking commercial information. They're in the information-gathering business, after all. Let some of the information you seek pertain to me spending money wisely or in fun or gratifying ways.

"What if I could go to my local paper to buy a house? I'm always interested in buying real estate," Dave writes. "If they sold me a house, then they would make money from the sale. A lot more than a few cents they make off me every year for the ads I ignore," he writes. "Maybe not a house. How about Internet connectivity? Or a movie date. Someone interesting to go to a baseball game with. These are things I pay money for. I pay a lot of money to go to games. How much I enjoy it is directly proportional to who I go with. All these things involve connecting people with people. So much money to be made there. Why doesn't the news industry help me meet interesting people? Maybe that's why Facebook makes so much money. Just sayin.'

"I'm also always in the market for better Internet connectivity. Could The New York Times help me there? We all live in the same city. They help me find good restaurants. Maybe if they helped me find better Internet; or, if they can't because it doesn't exist, if they helped to bring us better Internet by constantly beating the drum for it, which is something they can do and seem to like doing, that would be worth paying for. Beat the drum for new commerce, and then make it possible to buy the thing through your site."

He finishes: "There are honest ways to make huge money on the Internet. I think the message you're getting from your readers is that advertising is dishonest. The ads you show us net-net are junk. Jokes. Sad. Please stop this. Maybe the more distilled message is this: Stop talking so much. Listen." So that's Dave, posted a few days into this whole adblocking business.

And I read that, and I had this question, which I've never seen, at least in this context, asked before, which is how much stuff should there be on the Internet? That is, how big should the Internet be?

Leo: As big as it - it doesn't - there's no constraint on it.

Steve: Well, how many people should it be able to support?

Leo: Ah.

Steve: What I'm - yeah. What I'm getting at is that, because of this advertising model, which basically is about creating pages that seem more and more to be for the purpose of being viewed, rather than like having a real reason to exist, what's happened is - in fact, I think I was at Re/code the other day, and for some reason all of the links were

interesting. And I found myself thinking, wow, when have I ever found, like, a lot of stuff at a site useful or interesting?

Generally I use Google as the index, and I go somewhere for a particular purpose, for one thing. I ignore the ads. I read what I'm there for. And then I hit "back" in order to check out the next finding that Google had for whatever it is I was looking for. So I guess my point is that the Internet has become full of junk. I mean, there's just a bunch of crap on the Internet. And we've all kind of gotten used to it.

Leo: Well, no, wait a minute. Wait a minute, Steve. One man's crap is another man's treasure. You found Re/code useful. I guarantee there are lots of people who would have said, "There was not one link there I found useful."

Steve: Okay.

Leo: The Internet could be as multivariate as there are people. And I don't know if you'd get a general definition of crap. In fact, at this point, the only thing everybody agrees on is that ads are crap.

Steve: So wouldn't the prevalence of advertising tend to inflate the number of pages on the 'Net for their own sake? We've talked about, for example, how sites are now breaking articles up and making you go to successive pages...

Leo: Yeah, that's a perfect example of your premise, yeah.

Steve: Yeah. So it seems to me that, you know, we're going to go through some sort of an upheaval. Dave is sort of questioning other models. Essentially, because we've had advertising, and because it's been the classic no-brainer, you create pages, you put ads on the pages, you generate - now, yes, the more popular your pages are, the more popular your site is, the more traffic you will get. And so the revenue you generate from the ads on the pages increases. So if you bring more people to your site in this contest for eyeballs, then you're going to get a bigger share.

Leo: You're reinventing from first principles something that's already happened and died, which was called "demand media."

Steve: Ah.

Leo: So somebody said, gosh, you know, really the best way to make money on the Internet would be to figure out what people are searching for, create short, heavily advertise-laden pages that answer that specific query, and demand media was born. And so, for instance, there'd be a page on rattlesnake belt buckles with fairly mediocre, or in some cases terrible, content because one thing you don't want to do is spend money on this content. You want to generate a lot of it because you have to cover everything. What you do is you look at what searches are going on, and then

you create a page, and you load that page with ads. And everybody thought this was going to be the next big thing.

Steve: Gosh. That's terrible.

Leo: It died. Well, it died for a couple of reasons.

Steve: It deserved to die.

Leo: It deserved to die for a couple of reasons. One, because Google deprecated it. So Google has, in its page rank, really a way of measuring quality. And, you know, kind of an automated way.

Steve: Well, remember the inbound links model was the thing that originally made them famous.

Leo: Works great.

Steve: Who's going to link to that nonsense page?

Leo: Exactly. So what Google did is they started noticing what's a demand media page, and they just lowered it in the search results, and that effectively put those companies out of business. So we've gone through this cycle. But I think that what that points out is that there is a kind of a natural push towards better content. Just doing what you suggest doesn't actually work because it's been tried.

Steve: Well, I guess I was - I didn't mean to be suggesting something.

Leo: Well, I mean, you're saying there's a natural tendency, the way the Internet is constructed, to create this kind of motion towards just the content doesn't matter, get the ad views, make as much money as you can.

Steve: Ah, okay, right, right.

Leo: And in fact, it turns out, thanks probably to the intervention of Google, but it turns out that, yes, one would think that. But in fact there was a counter to that. And I think that the true counter is that, I mean, how many people saw those pages? It only made sense if you paid nothing, like literally a buck for somebody to write that page.

Steve: Right.

Leo: And so there were - but nobody really saw - and if Google starts to deprecate it, and the pages don't show up in search, then it doesn't work.

Steve: Well, and in fact we've also seen that on retired domain names. Domain names expire, and people grab them and put crap on them.

Leo: Right, link farms, yeah.

Steve: Yes, only because there will still be traffic going to recently retired domains from inbound links that don't know any better.

Leo: So we owe a real debt of gratitude to our hero Matt Cutts, who was the guy, you know, led the team at Google that did this, that basically tried to increase, make page rank work as it should, which is you get a higher ranking based on quality of content. That's the natural way that page rank would work.

Steve: You know, once again, you're a good example of doing it the right way. TWiT.tv exists as a domain and site for your enterprise. Much like GRC.com. You know, it exists as a container for all the stuff I've done. So many other sites are - they're nice, but - and the point was made, maybe it was yesterday on TWiT, that a lot of sites just sort of, they're here. But if users had any pushback, they would just go somewhere else. It's like, okay, I'm not going to unblock my adblocker. I'm just going to hit "back" and choose another link.

Leo: Well, then it's up to Google not to even make that link show up; right? I mean, I guess it comes down to how you get to a site. And if Google's search results favor better content, which I think they are trying to do, I mean, you might have a bad opinion of Google and think of other things that it might be trying to do.

Steve: No, they're my search engine.

Leo: That's why we use them; right? And in fact...

Steve: They find what I want.

Leo: ...they'd be out of business if they didn't do a good job of surfacing the best content on any given search. That's their job. And so I think that that's kind of handled. I mean, the fact remains there's stuff on the Internet that you, Steve Gibson, are not interested. But it's not Steve Gibson's Internet.

Steve: Right.

Leo: It's everybody's Internet.

Steve: Right. So I guess my point is that something is supporting - okay. So, okay, I know where we're kind of coming at odds. I'm suggesting that there's a lot of stuff that doesn't deserve to be on the 'Net which is only there because of ads. But if Google is doing its job, then no one would even know about those sites, and they wouldn't be generating any revenue, no matter what ads they had, because no one would ever go there.

Leo: Yeah.

Steve: Okay.

Leo: I mean, I think I come at it from a different angle. I think sites should respect their community, respect their readers.

Steve: Yes, yes. I think...

Leo: Respect their listeners, respect their viewers.

Steve: Yes. Yes, yes, yes.

Leo: That's what we do. And I think our success stems 100% from the constant, are we respecting our audience, are we building our community, are we doing what our community wants.

Steve: And you watch, you look at the feedback. I know you look at the feedback and are constantly tuning it in order to, you know, respecting your audience is listening to what feedback they offer and say, oh, you know, hey, there's a good idea. Let's do that.

Leo: And by pushing back against advertisers who want to do otherwise.

Steve: Right.

Leo: But we're lucky because I'm in the position where I do say no to advertisers. Now, maybe I'm reducing the amount of money I make. I mean, I think of Rene Ritchie, who works for a company called Mobile Nations. His site, and some of the other Mobile Nations sites, have a lot of ads. One of the Mobile Nations sites, I think it's Android Central, I can't use on mobile. It just doesn't work on mobile because it's crapped up. Is that in their interest? I don't think so.

Steve: No. I just again, my term is "perverse incentives," where if you put more advertising, for a while you get more revenue.

Leo: Yes. I understand.

Steve: And then the links begin to evaporate.

Leo: I guess that's my point is that the incentives really aren't really perverse. I think that they're - if Google does its right job. The problem is click bait. It's a game; right? So they come up with new ways to trick Google. A lot of this is gaming Google.

Steve: Right, the whole SEO evolution.

Leo: Yeah. And I feel bad in some ways for people like Matt because it's a full-time job to get around these games. Look what they did to Digg. Same thing happened to Digg.

Steve: Right.

Leo: If you think of Digg as a website designed to surface good content based on user votes, brilliant idea. Went under because these crap websites gamed it, and they couldn't figure out a way to stop it.

Steve: Yeah.

Leo: That's where you've got a problem.

Steve: So I already mentioned FireEye's detection. It was on the Forbes website from September 8th through September 15th. The Forbes.com website was serving content from a third-party advertising service that had been manipulated to redirect viewers to the Neutrino and Angler exploit kits. The FireEye guys notified Forbes, who worked quickly to correct the issue.

And I thought it was interesting, you know, the URLs where these ads were seen on the Forbes site, so it was Forbes.com, there was one, sabbatical-leave-work-leadership-careers-advice; should-the-fda-require-cv-outcome-studies-for-diabetes-drugs-before-approval, I'm sure CV is cardiovascular outcome. Then, under business was the-worlds-100-highest-paid-athletes. Under investing was the-grateful-graduates-index-2015-the-top-50-roi-colleges. So of course that's return on your investment for the cost of your education. And then under lists on Forbes was the-richest-person-in-every-state. So those pages, if you were unlucky enough to go September 8th through the 15th.

Leo: That's link bait. That's link bait, man.

Steve: Yeah. And, well, yeah, talk about the quality of the content. And you had a browser that was, like, standard configuration that was missing some updates. It redirected the browser through a chain of domains, as we've talked about before, and running one, two, three, four, five, six - I just counted them - seven different exploits contained in one Flash file. I have the link in the show notes to their explanation that shows the JavaScript, which is a simple little piece of JavaScript which is running a standard Flash movie, and downloads it, and tries to get into your system.

So they conclude, saying "Malvertising continues to be an attack vector of choice for criminals making use of exploit kits. By abusing ad platforms, particularly ad platforms that enable real-time bidding, which we've covered before here," they say, "attackers can selectively target where the malicious content gets displayed. When these ads are served by mainstream websites [like Forbes.com], the potential for mass infection increases significantly, leaving users and enterprises at risk."

So I looked at Crystal; Purify, which its full name is Purify Blocker; one called Blockr without an "E," B-L-O-C-K-R; Peace, which was Marco's and has subsequently been withdrawn, so it's only of interest if you happen to have it, you happened to buy it for $3 and didn't get your refund and still have it; something called Silentium; and 1Blocker. And as I mentioned before, my favorite for - I sort of came up with basically three. One, you know, sort of Mama Bear, Papa Bear, and Baby Bear.

For the simplest to use, not getting in your way, trivial one, essentially, is Chris Aljoudi's. He's the guy who took over the uBlock project. And of course, as we know, then that got forked to uBlock Origin. So I like Purify for its simplicity. It's $4, which is at the - it's the highest priced of any of them, although the power tool one is $3. So they're all in that, as I said, $1 to $4. And I don't think that ought to be a big hurdle. If you've found the right one, and you're using it, and it's saving you bandwidth and speeding up your page loads, then a couple dollars is probably not a big deal. It offers the very super simple whitelisting.

One of the tricks that they have all adopted is of using their little, essentially sort of a property page, where you're looking at a page, and Safari has that little, you know, "send this page somewhere." And that pulls up a second dialogue. You're able to go to the three dots at the end and then turn on - you'll see the name of the privacy blocker. Turn that on, allowing it to display there. And in fact you're able to, like, drag it up to the front of the list, if you wanted, so it'll be immediately there when you click on the "send this page somewhere." Tapping that icon for the blocker brings up its little settings sheet, which is where they typically put their whitelisting.

So again, if you just want to see the page you're on with no blocking, as I mentioned, you can hold down the refresh button until a menu pops up, and reload the page without any content filtering. Or, if you want to whitelist the site from then and forever, then you would tap the "send this page somewhere" and then get to the controls for that privacy utility and just say I want to enable whitelisting. So Purify was my choice for the simplest one. It offers that basically he's just taking all the responsibility for blocking stuff, using the lists that uBlock uses, which is sort of the same ones that uBlock Origin uses. So it's sort of from the same family and a nice simple blocker.

At the complete other end of the spectrum, the power user blocker, also in the Lifehacker benchmark, was the fastest of all of them, is called 1Blocker, the numeral "1" B-L-O-C-K-

E-R. And although you don't have to drill down, you really can. In the configuration of it you have 2,922 adblock rules, and you can individually see them. You can search them. You can put in DoubleClick, you can put in Google. So you could search these block lists. Nearly 4,000, 3,993 tracker blocking rules. You can block Twitter widgets with seven rules, Facebook widgets with nine rules, 18 other share widgets, web fonts, Discus comments, or adult sites. And in every case you're able to granularly select which of those as a class or individual rules.

You can also go to my.1blocker.com, which takes you to a beautiful web-based rule creator, where you're able to design your own rules and rule categories. And then he has a way of allowing you to import them into 1Blocker. It's free to download, so there's no charge to get it. But you can only enable one class of blocking, and everyone's going to want to turn on adblocking and tracker blocking and maybe Twitter and Facebook and share widgets and things. So immediately you get hit with I want $3. But again, $3 for this, if you really want a power user's content filtering tool, this is the one, 1Blocker. And we see from Lifehacker's benchmark of about 15, it was way ahead, I mean, substantially faster in their testing of eight different sites where they did URLs and measured the speed, the best speed improvement of all.

And for the in-between one, it's the least expensive, just $1, it's sort of the compromise between those two, is Blockr with no "E," B-L-O-C-K-R. It does offer a popup whitelist, which I think any workable blocker is going to need. And then you can choose when you pop up what you want to change. It will allow you some - so you have granularity, which is, for example, missing from Purify, the super simple one. This one, for only a dollar, you can choose ads, media, privacy, social buttons, and cookies. And so you can add exceptions for any site in a whitelist and also determine whether you want those blocked or not by default.

So that's the take. I'm sure we're going to, in fact, Lifehacker found a bunch more than I have found here. These are the first ones that I saw came out. If anything significant happens in the category, we will keep an eye out for it. And of course I'll let you know. But for now, I think on the low end, Purify, on the high end, 1Blocker. For those you pay either $3 or $4. Or if you want something simple but some control, some granularity, then it's Blockr, B-L-O-C-K-R, for a buck.

Leo: Turns out that the - so one of the reasons Marco pulled out is because his blocker, Peace, blocked the ads that were on his site, as well as his friend John Gruber's Daring Firewall site.

Steve: Yes, yes.

Leo: And there was no easy way, I mean, he didn't want to modify the Ghostery database. He thought that would be unfair.

Steve: So his ads were in the Ghostery database.

Leo: Yeah.

Steve: Right.

Leo: And they shouldn't be. And this is the real problem I have with adblockers. They're really a blunt sword.

Steve: Yes.

Leo: Because the ads on Marco's site and Gruber's site and Jim Dalrymple's The Loop are very unobtrusive. They're from a server called The Deck. They don't have Flash. They're aimed at Mac users because those all three are Mac-focused sites. They're as good as an ad can be. They look like a first-party ad. And they don't track. They're blocked by every one, including, by the way, 1Blocker.

Steve: I was going to say, the good news, though, is...

Leo: You could modify it.

Steve: Yes. With 1Blocker you could just put a couple characters of the domain name in. It will find the rule for you, and there's a simple switch for turning it off.

Leo: Yeah. In this...

Steve: I know. I know. I mean...

Leo: Okay. Let me show you the URL filter that you have to add. It's basically grep. You know, it's regular expression. I have to say that's not going to happen. Nobody's going to do that. And that's the real problem with these blockers is you're punishing everybody who has ads of any kind. And it's really the small custom sites, run by one or two people, like the ones we just mentioned, that are going to suffer the most.

Steve: I think, as we've said, this is a move to another model. And so this is the users of the Internet saying this got out of control. We're saying no. And in a sense it's, you know, it's an overreaction. It is blunt. And unfortunately it hurts the good ads as much as the really obnoxious ones.

Leo: The responsible advertisers and sites that respect their users. I'm just glad that TWiT is not a blog, is not a web-based enterprise because we'd be out of business.

Steve: Yeah.

Leo: And I just - really sad. I don't, I just don't - I completely understand, completely understand the very legitimate reasons for doing it. I wish there were a

tool - maybe that's what Adblock was trying to do with their, what was it, approved ads system?

Steve: Well, and you heard that they contacted the various iOS blockers and tried to sell them on permitting...

Leo: That's not what people want.

Steve: No.

Leo: People don't want ads anymore. But what they don't understand is...

Steve: People want to protest.

Leo: ...that is, as Clay Shirky, I mean Seth Godin says, this is a hundred-year-old model, free ad-supported media that trades content for your attention. And you're breaking it. And, you know, it's not completely - it's not the user's fault. I understand. It's probably the advertisers who broke it. But, boy, you're going to lose a lot of great stuff. I'm sad. And what you're going to get is the sites that get around it. The crap stuff is going to - this is going to increase crap, not decrease it, I bet you. But we'll see. In a few years we'll know. I'm just glad I didn't do a blog.

Steve: It's funny because a couple years ago Mark Thompson was saying, Gibson, you don't have any ads on your site. You know, it looks kind of strange not to have ads on your site.

Leo: Yeah. You need more ads.

Steve: Ads are supposed to have ads. And I said, oh, yeah, well, I'll get around to it someday. But now I don't think so.

Leo: No, I don't think so.

Steve: No.

Leo: You have a better model. You sell SpinRite, world's best hard drive recovery and maintenance utility, GRC.com. Just go there, you buy it, that supports Steve. There's lots of free stuff there. Really that's one of the things that's great about that site is all the free stuff, including a feedback form, if you want to ask questions for next week.

Steve: Yup.

Leo: GRC.com/feedback. He's got 16Kb ad-free versions of the show over there, as well. If you can stand the low quality, you won't hear any ads. That's the upside. He also has 64Kb audio and full transcriptions, which is great. And we really appreciate you doing that. We have the rest. We have the high-quality - everything we have has ads in it, I'm sorry, 64Kb versions of the audio as well as three different versions of video at TWiT.tv/sn. You'll also find us in most of the podcast applications, in fact all. I don't think there's anything that doesn't have Security Now!. Find it and subscribe, and that way you won't miss an episode. You think Q&A next week?

Steve: Yup, let's go for a Q&A next week.

Leo: All right.

Steve: So send your questions in. I will go through them. I did go through them for this week because this was nominally supposed to be Q&A. Everybody was wanting to talk about ads and adblocking, so I thought, well, okay, we'll talk more.

Leo: No, I'm really glad you did. Yeah, I'm really glad you did. It's very timely. All right, Steve. We'll see you next time on Security Now!. Bye-bye.

Steve: Thanks, buddy.