Mozilla's Tracking Protection

Description: Leo and I discuss the week's most interesting recent security events and a bit of miscellany. Then we examine the revelations about the current state of Internet user tracking arising from Mozilla's Firefox tracking protection instrumentation.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-512.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-512-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Yes, I know you want to know about the LastPass hack and what it means to you. That's coming up. We'll also talk about that horrific hack at the Office of Personnel Management. And Steve will talk about a very controversial new switch in Firefox that turns on tracking protection. Good or bad, Steve and I will debate, next.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 2^9, recorded Tuesday, 6/16/2015: Mozilla's Tracking Protection.

It's time for Security Now!, the show that protects you and your loved ones online with the Protector in Chief himself. I raise a glass to Mr. Steven "Live Long and Prosper" Gibson.

Steve Gibson: Tiberius Gibson, yeah.

Leo: I can only do it with my right hand. He's doing the...

Steve: Well, you know, I've been rewatching the Jean-Luc episodes of Star Trek, just because...

Leo: I love him. He's good.

Steve: It was, I mean, it stood the test of time.

Leo: 'Deed it did. 'Deed it did.

Steve: They were a great seven seasons of - and we have nothing like that now. I'm just - I'm amazed that there just isn't anything fabulous. There's some horrorful-looking thing - "horrorful," that's a new word - on the Syfy channel coming up. I made a note of it in the show notes. But it's like, it's one of those low-budget, the people cannot act, and whoever wrote it shouldn't be writing, and, oh, boy. But it's a desert for those of us who love science fiction.

Leo: It is, indeed.

Steve: So all kinds of stuff. Everybody wants to know about sort of the inside, what does it really mean, this LastPass network breach that just yesterday hit the news. I thought it was interesting that Joe gave PCWorld an exclusive interview yesterday evening.

Leo: Joe, the CEO of LastPass.

Steve: Yeah, Siegrist, the chief bottle washer, the coder, the original coder and so forth.

Leo: I know. When I read "CEO," it feels like, oh, he's just a business guy. But, no, he's the guy.

Steve: Yeah. He's making the technical blog posts. And he's always been my contact person. When I first found LastPass, it was with him that I corresponded. And one of the things that impressed me so much was that he just told me everything I wanted to know, unlike so many other services that hide what they're doing; and, therefore, it's not that I mistrust them, it's that I can't trust them because they're not telling me how it works. And, I mean, so everything with LastPass from the beginning was here's what we're doing. He wasn't embarrassed or shy or ashamed to just say, this is what we're doing. And there was a sense of, if I have any ideas, or if there's anything wrong, we're going to fix it.

I mean, it's very much, well, it's exactly the approach I've taken with SQRL, where the whole protocol has been hashed out and thought through and pounded on in a public forum with a bunch of other smart guys because the goal is let's just get it right. So that's what really impressed me from the start was I was able to describe on this podcast to our listeners why I had chosen it as mine, how it was truly TNO technology, and how they made it as strong as it was clearly possible to do. So we'll talk about that.

We did get even more bad news from the Office of Personnel Management. We discussed the bad 4.1 million record or individual hack last week. Turns out there has been another one, and it's much worse, and much bigger. Then there's this question, I guess it was a newspaper in the U.K. claimed that China and Russia had obtained and decrypted Snowden's entire document store and that, as a consequence, people were being pulled from the field. So we'll take a look at that. I saw with a little tear that our browsers may be losing native FTP. So for some set of users, that'll be interesting.

There was a mistake made with border gateway protocol, a big one a couple days ago that affected the whole Internet, I want to cover briefly. Some Wikipedia news. A really interesting reliability report about SSDs that was performed by Carnegie Mellon and Facebook, using Facebook's massive datacenter servers and the nature of SSD failure. We learned a lot from that. And then the main topic is I listened to you, gosh, I guess it was maybe last week's This Week in Google?

Leo: Oh, yeah, about Safari ad blocking?

Steve: A lot about - you and Jeff and the rest of your panel, really interesting discussion about ad blocking, the ethics and morality.

Leo: Mike Elgan was there, as well, yeah.

Steve: Right, right, right, right. I mean, and great comments, I thought, from everyone. But from a technical standpoint there were some things that I wanted to add. And I already had on my list, because I'd been mentioning it for a couple weeks, that Mozilla, back with Firefox 35, they added something called tracking protection that's been on my list to talk about. And I thought, okay, now is the perfect time because, thanks to the instrumentation they have in Firefox, they've been able to learn a great deal about what websites are doing. And so I just want to have a discussion about that issue a little bit, but also from a technology standpoint, the degree to which I would argue, first of all, the incentives are wrong, that is, they're perverse incentives in the industry as it is now. And that it's beginning to get a little out of control. So I think a great podcast for everyone.

Leo: So excited.

Steve: And on the first page of my show notes, of course, I try to do something that's fun or interesting. Last week was a screenshot of all the guys from "Silicon Valley." And of course this was on the following Sunday. Last Sunday was the second season finale.

Leo: Did you not love it? Was it not great?

Steve: Oh, yes. This year, this year was just so fun. I mean, and the idea that this guy is dying on a mountainside on a webcam, and they could care less.

Leo: [Crosstalk] says, "Oh, we are so lucky he fell. It's such a great test."

Steve: And then he announced he was going to have to start drinking his own urine. And they said, "Oh, my god, this is going to drive our traffic."

Leo: Fantastic news. Oh, I hate to even laugh at that. So you put an image of, well, let's see, shall we zoom in on it so - it's a cartoon; isn't it?

Steve: It's a cartoon. It's a great computer-based cartoon showing a boxing ring. And the announcer's holding a microphone saying, "And in this corner, we have firewalls, encryption, antivirus software, et cetera. And in this corner, we have Dave." And there's this goofy-looking Dave with a T-shirt that says "Human Error" on it.

Leo: Human error.

Steve: So it's like, yes, despite all of our best efforts.

Leo: Not much you can do.

Steve: And in fact, reading between the lines of what Joe said to PCWorld about what they found on their network, I'm thinking it was some guy, someone there got - an individual's machine was the entry point because he talked about some anomalous traffic on the network when no one was working. Well, okay, that connection says that it would be traffic driven by someone working, meaning, like, someone's workstation. So anyway, we'll talk about that in a second. Oh, and I did want to also say Merlin Mann is a fabulous guest.

Leo: We love him.

Steve: I mean, just whenever you can drag him out of hiding. I remember the evening you and I and Amber and he spent together.

Leo: Oh, in Toronto, yeah.

Steve: In Toronto.

Leo: I love Merlin, yeah.

Steve: Just hanging out for a nice evening. He really brought a lot to the show, I thought.

Leo: Agree, agree, agree.

Steve: So I just wanted to say, you know, Sunday was a great show.

Leo: Good, thank you.

Steve: Okay. So LastPass's network breach. Of course the press went crazy, and my Twitter feed was literally useless yesterday because, I mean, I thank everyone for

making sure that I knew.

Leo: Steve, did you know? Had you heard? Did you know? There was a problem. Did you hear that?

Steve: Yeah. And in fact what Joe said to PCWorld was they dramatically underestimated the media's reaction to his blog posting and the email they sent out. Turns out that they have - LastPass has grown so popular, and they had so many people to notify, that it became very difficult to get the mail out.

Leo: Oh, dear. Oh, dear.

Steve: And so they're now - they used an Amazon service, and they're going to use Amazon's scalability moving forward. And I can attest to the fact that people went crazy because I was unable to change my master password, as the good advice was, until later in the evening, when the initial wave of just everyone attacking the server hit. And in fact another thing that's interesting is that, as a consequence of their security being so good, which I'll describe in a second, but there's one point where they do a hundred thousand iteration PBKDF2, Password-Based Key Derivation Function.

Leo: Isn't that amazing?

Steve: Well, yes, except that that means there's substantial overhead for them any time a user changes their password.

Leo: Which means, uh-oh, everybody's changing their password.

Steve: Yes, which means, I mean, so they did this deliberately to strengthen against exactly the attack that they are worried may have happened. And it's necessary, I'm being very careful with the way I phrase this because they don't actually know there was exfiltration. They don't have evidence or knowledge or logs.

Leo: That's interesting. They're just assuming it because that's prudent.

Steve: Yes. They're being that overly cautious. So again, this is another reason why...

Leo: We love them.

Steve: ...they're a model for the way you do this. And they said, we routinely survey our network because that's the way you do it these days is you look for anything that seems suspicious, and then you go figure out what it is. So something seemed suspicious, some traffic that they didn't expect should be there. And so that's why I'm guessing - and specifically what Joe said was "when no one was working." So it was like, you know, after

hours. Workstations should have been idle. But if someone's workstation...

Leo: Ah. That's a giveaway.

Steve: Yeah, I sort of think so. So that was probably the port of entry, so to speak. And so something was there. Then they looked at what the traffic, where the traffic was. And they already have a very segmented system. So, for example, they were absolutely able to determine that none of the data, none of the bulk-encrypted data, which is the reason we have them, for synchronizing through the cloud our password databases, none of that was ever at risk. But where the anomalous traffic was on their network could have exposed a subset of what they're keeping. And, unfortunately, they are keeping secrets.

So, for example, the password side stuff was where this anomalous traffic appeared, which were the email addresses, the password reminders, the per-user salts, and the authentication hashes. So let's remember the way this system works. The reason I like it so much is that the user's email address, after case is removed - so it's case-insensitive because email is case-insensitive, and they don't want to confuse their hashing because they combine the email address and the user's password, and then they hash it iteratively. And I don't remember now what the default is because I had cranked mine way up. Advice on this podcast a year or two ago was, when they added - okay, now, you've got it set to 100,000.

Leo: The default is 5,000.

Steve: Okay.

Leo: Which is probably enough. I mean, they were doing 100,000. What do you recommend?

Steve: Okay, so what I recommend is a number with no zeroes.

Leo: Oh.

Steve: So everyone should choose - and you don't want to show it to us, Leo, because that would be a nice thing for no bad guys to know.

Leo: Why would they want to know that?

Steve: Because that really foils them. I already had, like you, a five-digit number. But it was completely random, a random five digits. And they caution not going higher because this is done browser-side.

Leo: They say don't go below 20,000.

Steve: Good.

Leo: It says it's recommended that you keep the number - this is the message I'm getting from LastPass when I change it. It's recommended you keep the number of password iterations, oh, below 20,000. Otherwise you may experience a significant delay. And of course that's the point of all this extra computation, right, is to slow down brute-force attacks.

Steve: Correct. And if the bad guys assume that the users have followed that advice, then they won't assume that we're using a larger number. And they will also assume that we chose something with a lot of zeroes. My point is it matters what you choose here. It should be random, and it should have five digits, and the first digit should be greater than two. So what happens is your email address and your passphrase are essentially hashed iteratively that number of times, that crazy number of times.

Now, none of that may actually matter, that is, the secrecy of that. I shot a note off to Joe, but he's just been underwater since this happened, because there were a couple of things I needed to find out. So it may well be that that number is part of the data that they are keeping and would have been attacked anyway. So keeping it a secret may not matter.

But the point is we're doing - one of the reasons I immediately knew this was the right architecture was that our browser does that locally. And the result of that hash is our identity. So it's an anonymous identity which we are providing to them. And then it's hashed again to provide the key. I think that's the order. It's been a long time since I've thought about this. But that provides the key for our decrypting the data that our browser stores locally, so that it's a very strong local store. Then they get that thing, that blob from us, which has already had the crap hashed out of it, and they do it another 100,000 times because they've got big iron servers. And, I mean, and they want to take responsibility.

Oh, and also remember that I don't think that local hashing was - it either wasn't there, or it wasn't as strong. I don't know if they - they may have had a fixed number in the beginning, in order to produce a strong hash. Or maybe it was a smaller number. I don't quite remember the history of how this evolved. But they're accepting something that they want to protect. And they figure, hey, we've got big, strong servers. We're going to further hash it. So they take what they get from us and do it another hundred thousand times. Okay. But the point of this is, what may have been exfiltrated is all of that data. The email address, which they do have in the clear, they have that because that's how they send us notices.

Leo: They have to, right.

Steve: Yes. And so the email address is half of the secret, which is not a secret. Obviously it's known. Then the passphrase is what's added to that. And then it goes through all this hashing. The point is bad guys could perform a targeted attack. The other thing that they've done right is they have what's called a "per-user salt." And many of our longtime listeners know all about what that is. The idea is you don't do the same algorithm, exactly the same algorithm, to every account on a big server because that allows you to create tables, the famous rainbow tables, which could be used as lookups, where you would only have to do the computation of, like, a hundred thousand hashes

once for all possible passphrases. And that would give you the result, which then you look for the result in the database, and then that tells you what the passphrase was.

Instead, every single one of these uses a random, it's called a "salt," which is mixed in for that one account, which means there's no way to do anything for everyone. The only attack then would be for a bad guy to take the email address, which assuming that anything got out at all, to take the email address of the account - and notice that, I mean, there's some information there. That tells them probably who that account is. And there are no doubt many powerful and famous people who are using LastPass. So if this got out, the record is identifiable by that high-value person's email address. So that tells them who's worth attacking. So they would then have to take that email address and that user's salt, that user's account salt, and start making guesses of what their password may be. Now, here is why a strong password is important. Hopefully, you know, this isn't Paris Hilton, and she's used her dog's name again, because that would be very bad.

Leo: That would be wrong.

Steve: Yes. They would emulate what the browser does and make - if the number of browser iterations is part of what got taken, then they would know what the browser is doing. But so they would emulate what the browser does. Then they would emulate what LastPass's servers do, that is, the 100,000 additional PBKDF2s using SHA-256. Now, SHA-256 is, unfortunately, the bitcoin hash. So we now have incredibly mature, ultra high-speed SHA-256 hashing ASIC hardware, which is in the terahashes per second range. I mean, it screams. So this is acceleratable to a great degree. And if the guess for the password was right, then they would be able to determine that.

Okay, now, LastPass has protected us such that you cannot log in from a new device or IP as of before the announcement, unless you do an email confirmation. So they would not have compromised this person's email also, almost certainly. So that wouldn't help. That would probably not help them. And we should all have just changed our master LastPass password so that, if the worst happened, and data actually escaped, and an individual was targeted, and that individual had either a weak password or just brute-forcing finally found it, then it wouldn't help them anyway because you would have changed your password for LastPass, and they're out of luck.

But if you used the same master password anywhere else, that is, if your LastPass master password was something you reused on other sites, this breach could theoretically give bad guys a way of obtaining the password you used for LastPass and knowing who you are. They would have your email address and could potentially use that to log into other sites. Which is why the advice in Joe's blog entry and in the letter that went out was, if you used your LastPass master password somewhere else, you really need to change those sites because, again, so everyone should now have a good sense for the tremendous amount of effort that stands between bad guys and actually getting anything useful, if in fact there was a breach that did exfiltrate this information. We have to assume they had reasonable cause to believe it.

I don't know, we don't know in detail what evidence LastPass has because this has obviously been a big inconvenience. There's been some reputation cost to them. And we know that, I mean, their servers were completely, virtually offline, essentially, under the load, while everyone was frantically trying to change their master password. Although there really wasn't any hurry.

Now, I've seen some online, actually, Simon Zerafa tweeted me a link to some of the work that had been done over in the high-speed hashing projects. And I saw a number that looked like maybe 8,000 guesses per second, with strong GPU-based technology. Now, again, ASICs have blown GPUs away in terms of SHA-256 hashes. So even that is low. But for what it's worth, the only attack that is feasible is a targeted attack based on what your email address for LastPass is. So if that was anonymized or a separate Gmail account or something that doesn't look tasty, then you probably, you know, it'd be unlikely you'd be a target because there's no way to do a mass crack of this, thanks to the per-account salting.

So again, we know that it is incredibly difficult to defend a contemporary high-value network against bad guys. It's hard to imagine a higher value network today than LastPass because it's known that they literally have the keys to the kingdom. They have deeply encrypted, well protected, but still there, everybody's cloud-synced master cross-website Internet identity, essentially, everybody's usernames and passwords that are being used to log in. So it creates an incredible target. I mean, incredibly valuable.

And I can't resist the opportunity to mention that this is another one of the many things that completely disappears with SQRL. There is no cloud-based database in the sky. There's no need to synchronize independent copies of password databases because there's no more password databases. And in fact this is a complete SQRL identity. That's a SQRL identity. And all clients are able to use that QR code in order to clone that identity between SQRL clients. So my client for Windows can display that, and you just snap it, you scan it with your phone, and now your phone has your identity. And they're automatically all synchronized among all your devices and all the websites you ever go to, and there's nothing for anyone to steal.

Leo: Okay. But just to be practical, it's not available now, and nobody's using it now, and who knows how widespread SQRL will be. So let's give some practical information here. What do you recommend for people to do who are LastPass subscribers? Should they change their email?

Steve: No. Okay. So the point is, only by being a target of an attack are you in danger. So you want to do two things. You absolutely need to change your LastPass master password because, if this data got out, if you were targeted, if you had a weak password, and if they could figure out how many iterations to use, then they would be able to confirm the password that you were using on LastPass and have your email address, therefore being able to impersonate you to LastPass and get all of your LastPass login authentication. So you absolutely want to change your master LastPass password. The second thing...

Leo: Right, and I did that. And of course I promptly forgot what I changed it to, but that's another matter entirely.

Steve: And, you know, that happened to many people.

Leo: It's really easy to do, especially because you're using a long - the real challenge is you want to come up with a really nice random long password, but you also want to be able to kind of recreate it. And I use poetry, the first letters of lines

of poetry. But I misremembered the poem, apparently, so now I have very - but that's okay. Others have had this problem, too, I'm sure.

Steve: For what it's worth, many people had the problem. There are instructions at LastPass for reverting to your previous forgotten password if that happens.

Leo: Is that a safe thing to do? I'm surprised he even offers that, to be honest.

Steve: I am, too. But I think, unfortunately, I mean, these are all the problems that we have. Last week, when I demonstrated clicking the QR code, someone in the chatroom said, "How is that better than the browser memorizing my password?" And it's like, oh, my lord. What we're doing is we're shoring up an existing horribly broken system with one attempted solution on top of another. And so this whole issue, you know, we've gotten to the point now where we have to use different passwords for every site. So no one can remember those, so we have to use a password manager. But we're also using different computers and different devices, so we need synchronization. So that means that synchronized database needs to go into the cloud. And now, you know...

Leo: That's the problem, by the way. If you didn't need synchronization, you could just have an encrypted file on your desktop. If you didn't care if anybody else saw it, you wouldn't need to do that.

Steve: Right, right.

Leo: But synchronization is what we want, and that's the problem.

Steve: Yes, and how many times have we been saved by LastPass keeping things synchronized? Because we do change a password over on some device on some site, and it's like, it's almost a miracle. I mean, it's wonderful that then you later go to a different computer, and it knows how to log you on there.

Leo: You know, I'm kind of lucky because I have two LastPass accounts. We have an enterprise account, and I have my personal account. And I had merged my personal account into the enterprise account, which you can do. You can, for ease of use, if you have an enterprise account, you can have both your LastPass accounts. And I had done that. So I changed the enterprise password, saved that, thank goodness, for my enterprise account, and everything's in there. So the fact that I can't get into my personal account's not really a problem anymore. Let me ask this. Is there a...

Steve: So my answer is, for your LastPass master password, it needs to be something big and random, and you have to write it down.

Leo: Put it somewhere.

Steve: And I don't mean it literally. You have to record it. So, like, do cut and copy and paste and put it somewhere else. Put it somewhere, maybe print it out, or maybe put it in some other place which you feel is secure.

Leo: The problem is, because of synchronization, you need it on all the other devices, too.

Steve: Yeah.

Leo: So you have the same problem. Is there a...

Steve: Or maybe put it in your wallet, but also make some change to it. Don't have it be the exact one. Leave out, you know, change a digit so that...

Leo: That's a good idea.

Steve: Yup. And when it doesn't work, then you'll remember, oh, that's right.

Leo: Oh, yeah, I've got to do this.

Steve: The last digit is actually - uh-huh.

Leo: Is there a better solution? Does this mean that we shouldn't be using LastPass?

Steve: No. I don't think anybody - okay. So in some subsequent postings, Joe acknowledged that some lessons were learned. And so, for example, one of the things they're going to do is they're going to rely more heavily on Amazon server scaling so that they're able to scale their performance to better deal with this kind of surge in everybody needing to do something massively computationally burdensome at once. And I got the sense that they're going to do some additional network segmenting. Ultimately, that's really what you want. We've talked about this, like in the Sony hack, where one of the horrible things about that was that apparently everything was available on a single network.

So where your network is about security, as is the case with LastPass, rather than about convenience, which was the case with a bunch of entertainment people in Hollywood at Sony Pictures, there it really makes sense to separate networks. It means that it's less convenient. But wow, you know, ultimately it's probably what you have to do when the reality is everybody makes a mistake, as that cartoon at the top of the show said. Despite your firewalls and your antivirus and everything, you've still got Dave, who is going to click on the link in the email from his mother and infect his computer, and he's in a trusted network in a company that's all about trust.

So my point is, to answer your question, I'm not leaving. I'm staying with them. I mean, I've always thought it was sort of odd that oftentimes employees are fired when they

make a mistake that taught them a lesson because you've just lost somebody who knows more than they knew before, and you're going to replace them with somebody who hasn't learned that lesson and hasn't been chastised. So I've never understood that. I mean, malice is one thing. Goodbye. But a mistake, you know, that happens. So all of their technology is bulletproof. There is, for example, there is nothing I can think of that they could have done more. Obviously, except not to have this happen. But given that it's happened, the environment that it has happened in is as secure as anything I could imagine.

And look at how diminishingly small the risk even now is. It is incredibly difficult for someone to obtain one person's password, maybe, which they've probably changed, and hopefully didn't use anywhere else because the whole point of LastPass is you don't have to use the same password everywhere. And LastPass gives you auditing facilities to look at all the passwords you're using that it's storing for you and verify that they're unique. So, I mean, it's not - maybe they were trying for more, and this is all they got, and they're not happy. I mean, maybe nothing's going to even be done with this. So again, from a practical standpoint, I can't think of anything more they could have done. Users should change their LastPass password and also change other sites' passwords if you were using them also with LastPass.

Leo: And never do that again.

Steve: Right.

Leo: You should also turn on two-factor. I mean, I didn't feel the need to change my master password. I did. But I feel like...

Steve: Oh, no, two-factor doesn't help you here.

Leo: It doesn't help you here.

Steve: You notice I didn't mention it at all.

Leo: You didn't.

Steve: No.

Leo: Okay. Tell me about that.

Steve: Two-factor is orthogonal to this. This is all separate. Two-factor comes in after you've authenticated yourself. So it would prevent a bad guy from logging in with your old LastPass password, if you did not change it. So that part is useful. But the two-factor still allows them to do this attack and determine what your LastPass password was. So if you did have second-factor, you're safe from them logging in as you. But you're not safe if you reused that password anywhere else.

Leo: Right. But you shouldn't have been reusing it. So I did not reuse it. So there is some value to having two-factor, if you don't reuse it.

Steve: Oh, absolutely, yeah.

Leo: Yeah. So nobody - the main point for me is I don't want anybody to get into my LastPass store. That secure store is the keys to the kingdom.

Steve: Correct.

Leo: Because everything's in there, including my social security numbers, my banking account, my banking account passwords, every - if you had my LastPass store, that would be terrible.

Steve: Yes. Well...

Leo: So, and there's no threat that that was compromised. That's encrypted with Trust No One encryption; right?

Steve: Well, and it's clear that they did do the network segmentation I talked about. They absolutely know that traffic on wherever they saw it meant A, but it did not mean B because they absolutely have these systems segmented. So that they've really done right. They were able to emphatically say that the actual databases were never exfiltrated.

Leo: Yeah. So to reiterate - oh, and by the way, could this have happened to KeePass, the open source password manager? It couldn't because that doesn't do syncing. So they don't have anything on a server anywhere.

Steve: Right.

Leo: So, but that's the disadvantage of - that's one of the reasons I don't use KeePass, because I want my LastPass passwords everywhere. And I use, as many of us do now, many, many different devices. So that's the convenience factor. So your advice is change your master password. And this is especially important if you've ever used, for reasons I wouldn't understand, your LastPass master password anywhere else. Change it. You should turn on two-factor authentication; right? Why not?

Steve: Yes. It provides additional protection.

Leo: And it's not a big onerous problem. And then go into your LastPass settings, and it's in the advanced settings. He's started to hide some of this stuff. He's rearranged his settings. So go into Settings. On the first page is an Advanced button. And you'll see...

Steve: Yeah. I couldn't find them. I was looking all over for it.

Leo: You have to search now, yeah.

Steve: And I had to go to Google. And I said, oh, down - and it's like, oh, there's an Advanced button. I didn't see it, yeah.

Leo: And then you go, you'll see iterations 5,000. That's the default. Bad for a couple of reasons. It's too low, and it's a known number. So you want to change that to a five-digit number, does not begin with two, and it should be random.

Steve: And you'll never be given a test on this, so you don't need to remember that. Just make something up.

Leo: Right. You never - you don't have to recover it. Again, it's just how many times it will rehash this. And you don't want it to be a number the bad guys might know. You want it to be a number they couldn't possibly guess.

Steve: Exactly.

Leo: And why not begin with a two? Because they're going to start - that gives them one less digit to guess because two is the limit, 20,000 is the limit.

Steve: Oh, just because the dialogue says don't go over 20,000, so they won't go over 20,000.

Leo: Start with two, yeah.

Steve: So we want to. Yeah.

Leo: Yeah. Oh, so do use more than five digits. Or use five digits and start above two.

Steve: Okay. So the only reason not to use a bigger number - bigger is better. But the warning that they say is it'll slow you down. Well, I have a - is it five or six? I'm not telling. But I've got a big number with crazy digits, and it wasn't slow. It had to go - I

watched the little bar.

Leo: Computers are fast.

Steve: Yes. I watched a little bar go across. And how often do you do that? It's not happening all the time.

Leo: And in fact, I have hundreds of passwords in LastPass. When I did that, it had to reencrypt every one of them. It takes 10 seconds.

Steve: Yes.

Leo: It doesn't take any time. Get a cup of coffee, it'll be over. All right. So a lot of the freakout is not necessary. But do the bright thing and change your master password.

Steve: Yeah.

Leo: And what is the minimum number of characters you should have? Fifteen? Twenty?

Steve: Yeah. Oh, yeah, I mean, again...

Leo: Make it long.

Steve: People are bad at generating entropy. So, I mean, it's why my crazy passwords page. It's like 25,000 uses a day. I think that's some scripts that are borrowing - that are getting entropy from GRC. But before that happened it was, like, 5,000 a day. I mean, people are using it all the time, just to get gibberish, because they like my gibberish rather than just gibberish anybody else has or could make up. So, I mean, I use it myself. I go to there, I copy something out of it, and then I mix it around, then I use it. I just use my own gibberish when I'm needing a password. And that's what my LastPass password is. I have no idea what it is. It's complete gibberish. But it's been recorded. So I really think that's what you want to do. I mean, it's inconvenient. But while we're stuck with passwords, I think LastPass...

Leo: That's the inconvenience.

Steve: Yeah.

Leo: Is this crappy system.

Steve: Yeah. And it's belt and suspenders. We're reacting to the nature of past attacks, that databases are lost on websites, so we have to use different passwords on each site, and that attacks, that they're brute-forcing them, so now we need to use passwords we can't remember. And so if we're going to have gibberish spread all over the Internet, we need something to keep it.

Leo: Hold the gibberish.

Steve: And we all have multiple devices, so now we need them synchronized. So now we need a cloud.

Leo: By the way, somebody's asking in the chatroom, this will never go away, thank you xkcd. I love xkcd. But they were wrong when they asserted that a passphrase with English language words is better than a random password. Make a random long password. A passphrase is not better.

Steve: Right. And was it Bruce Schneier, somebody who agrees with me said, "Make it long. Make it random. Write it down." We know how to manage little bits of paper. So manage the little - and then, of course, and the famous comeback is, oh, yeah, the Post-it note under the keyboard. Okay, well, don't leave it under the keyboard. Put it in your wallet. And, as I said, make a change to it so that literally it doesn't work as is. And when you forget that, and you enter it, and it doesn't work, then you'll go, oh, that's right, I made a little change. And then make that change, and you're in. That's the only way to be safe.

Leo: I have a method.

Steve: I think that's good.

Leo: That I don't want to tell anybody.

Steve: You shouldn't.

Leo: But I have a seven or eight or nine-digit number that I remember, and I just append to that.

Steve: Yes.

Leo: To whatever. And that's never written down. But that's in my head. That's one way. Right?

Steve: Yeah.

Leo: And you've talked about padding. In fact, you have some great stuff on GRC.com about padding.

Steve: Yeah, haystacks.

Leo: Haystacks. All right. So I think we've covered this. You know what, I'm really glad that we did cover it, though, because there's a lot of misinformation and a lot of panicking, frankly.

Steve: Yeah. And I saw a lot of tweets from people who were saying, "Thank goodness Security Now! is today because we'll get the full readout about what this actually means." And so all of our listeners now know that, like, exactly what the risk, what the nature of the risk is and that I'm continuing to be a happy LastPass user. The nature of security is that mistakes are going to happen. And what you need is an architecture that does everything it can to minimize the damage from mistakes. And what we have just seen is the operation of exactly such an architecture.

Leo: Well done, as always, Mr. Gibson. Continue on. The world has been saved once again.

Steve: Speaking of disasters, the Office of Personnel Management...

Leo: Ohhh.

Steve: ...has even bigger troubles than we knew. The first breach from a week and a half ago was believed to be 4.1 million records of a certain class. Now we learn that there was a second intrusion involving many more pieces of even more sensitive data. I cut down an Associated Press article, and I'm paraphrasing it. But I just saved the juicy bits.

So the AP reported that hackers linked to China - and American officials have said that the cyber theft originated in China and that they suspect espionage by the Chinese government, which of course the government, the Chinese government has denied any involvement. So "Hackers linked to China gained access to the sensitive background information" - and I think you were talking about this on TWiT on Sunday, too, Leo, the sensitive background information, which is different than what was leaked before.

Leo: Oh, and worse, you may not even have had a government job, and they still have that information; right?

Steve: Correct, "submitted by intelligence and military personnel for security clearances, in a cyber breach of federal records dramatically worse than was first acknowledged. The forms, which authorities believe may have been stolen en masse, known as Standard Form 86" - I love that, 86. Of course that's standard jargon in the restaurant industry, too. It's like, oh, yeah, we 86'd that - "require applicants to fill out deeply personal information about their mental illnesses, drug and alcohol use, past arrests, and

bankruptcies. They also require the listing of contacts and relatives, potentially exposing any foreign relatives of U.S. intelligence employees to coercion. Both the applicant's Social Security number and that of his or her cohabitant, if any, is required. And beyond Social Security numbers, the data include military records; veterans' status information; addresses; birth dates; job and salary histories; health insurance, life insurance, pension information; age, gender, and race data."

So, I mean, basically stunning, comprehensive information about individuals. But the scope of this is equally disturbing. In a statement, the White House said that on June 8th investigators concluded there was, quote, "a high degree of confidence that systems containing information related to the background investigations of current, former, and prospective federal government employees and those for whom a federal background investigation was conducted" - and, by the way, this goes back to the 1980s. So this is old stuff, too. And as you said, Leo, even if you didn't get the job, they still held all this information - "may have been," said the White House, "exfiltrated." Joel Brenner, who's a former top U.S. counterintelligence official, said: "This tells the Chinese the identity of almost everyone who has a United States security clearance."

Leo: Oh, my god.

Steve: "That makes it very hard for any of those people to function as an intelligence officer."

Leo: And they're claiming Snowden got people in trouble. This is far worse.

Steve: I know. I know. In fact, I thought of this when I was - because we'll speak about the apparently bogus story...

Leo: No, it's a Rupert Murdoch slam piece. It has no reality.

Steve: Exactly, that China and Russia had acquired and decrypted his stash of documents.

Leo: Didn't need to. They already got everything.

Steve: Directly from the source.

Leo: Oh, my god.

Steve: It was, in fact, it was a lot fresher. Snowden's stuff is old now. That's old news. Now we've got the last four years' updated intelligence information, directly from the Office of Personnel Management. So the White House statement said - oh, anyway, so continuing Joel Brenner's comment, he said: "That makes it very hard for any of those people to function as an intelligence officer. The database also tells the Chinese an enormous amount of information about almost everyone with a security clearance. That's

a gold mine," says Joel. "It helps you approach and recruit spies." And of course there's been some concern that it could be used as blackmail material.

Leo: Horrible.

Steve: That is, you know, we know about you, so do this for us. Just this little thing, and we'll keep your secret. "The White House statement said the hack into the security clearance database was separate from the breach of federal personnel data previously announced" - and of course we know that - "a breach that is itself appearing far worse than at first believed. Nearly all of the millions of security clearance holders, including some CIA, NSA, and military special operations personnel, are potentially exposed in the security clearance breach, the officials said. More than 4 million people had been investigated for a security clearance as of October 2014, according to government records." Okay. But that was the newer take on the previous breach.

"But in this newly released hack of standard personnel records announced just last week, two people briefed on the investigation [that the AP is quoting] disclosed Friday that as many as 14 million current and former civilian U.S. government employees have had their information exposed to hackers." And these are the records that I talked about, all of that stuff going back to the 1980s.

Leo: Including Snowden's records, by the way.

Steve: Yeah.

Leo: Ironically.

Steve: "Since there are about 2.6 million executive branch civilians, the majority of those records exposed relate to former employees. Contractor information has also been stolen, officials said." So anyway, just wrapping up, they said: "The personnel records would provide a foreign government an extraordinary roadmap to blackmail, impersonate, or otherwise exploit federal employees in an effort to gain access to U.S. secrets, or entry into government computer networks." That's how you guess, you know, social hacking of all kinds. "Outside experts were pointing to the breaches as a blistering indictment of the U.S. government's ability to secure its own data two years after a National Security Agency contractor" - of course they're referring to Edward Snowden - "was able to steal tens of thousands of the agency's most sensitive documents."

Leo: Yeah, but he was in the building. These guys weren't even in the building. Terrible.

Steve: Yeah. "After the Snowden revelations about government surveillance, it became more difficult for the federal government to hire talented younger people into sensitive jobs..."

Leo: Oh, oh.

Steve: Uh-huh, "particularly at intelligence agencies," said Evan Lesser, managing director of ClearanceJobs.com, a recruiting firm.

Leo: Imagine how hard it's going to be now.

Steve: And ClearanceJobs.com matches security clearance holders to available slots. And they're saying they can't get young people because young people are just saying, uh, no thanks. And so he said: "Now, if you get a job with the government, your own personal information may not be secure. This is going to multiply the government's hiring problems many times."

Leo: Jiminy.

Steve: And then anyway, then of course Mike Masnick at Techdirt picked up on this and covered the story. And he ended his coverage by saying, "And yet, this is the same federal government telling us that it wants more access to everyone else's data to, quote, 'protect us' from cybersecurity threats, and that encryption is bad. Yikes." And remember that these are also the people, you know, Donna [Seymour], I forgot her name, who's the CIO, who has not been heard from, by the way, in the last couple weeks, that, yeah, encryption is new technology, and we're still working to deploy it.

Leo: Wow. Wow.

Steve: Mega.

Leo: Mega mega.

Steve: And so this is on the heels. We have in here, I did want to cover the question of, and you already know the answer, Leo, did China and Russia in fact obtain and decrypt Snowden's document cache. And Glenn Greenwald, who has a horse in this race, and no one would say that he's neutral, but he did just tear the reporting apart. One of the main factual pillars on which this stood was that Snowden had documents with him in Moscow, and that Greenwald's partner met Snowden there, and that there was some document exchange. None of that ever happened.

And so the story was just - the U.K. paper story was just laced with falsehoods and, of course, citing unnamed government employees. And of course the whole story was that this was going to put their agents at risk, but it didn't say who. Not that they would, but that they were pulling them out of the field because their lives were in danger. And so anyway, any fair reading of this looks like this was not the case.

And also, I mean, everything we know about Snowden is that, whether you agree or disagree with everything that he's done, he really appears to have never been anything

other than straightforward. And we know that he understands encryption. He stated that he destroyed his personal copies of this so that he could not, specifically so that he could not be coerced into giving it away. The cache is in the possession of the press. And he had none of it when he went to Hong Kong and then on to Russia. So I don't see any reason for that not to have been true. And you cannot disclose what you don't have. So this looks like just completely made up.

And old timers among us will maybe find this interesting. I did. On the bugzilla.mozilla.org site, Bug - they're calling it a bug - 1174462, titled "Remove built-in support for FTP."

Leo: Hmm.

Steve: Yeah.

Leo: We're old-timers, though. And first of all, FTP's not secure.

Steve: No. And in fact I use Wget now. I don't think - I can't remember the last time I actually..

Leo: Yeah, me, too, yeah. Curl or Wget.

Steve: Yeah. You just don't put ftp. I guess the only place it would hit it is if a really old, creaky site, like something that maybe Jerry Pournelle would have up...

Leo: I bet he still has an FTP page.

Steve: I bet he has FTP links. That's my point, is there would be links that would be ftp://.

Leo: Ftp://, yeah.

Steve: Where you would use an FTP protocol to download code, rather than HTTP.

Leo: The point is, anybody who's using FTP or SFTP is probably going to use an FTP client rather than the browser. And so that's just, yeah, I think it's probably a good idea to take it out.

Steve: And so in the Chromium, Chrome is doing it, too. So this was sort of where it came from. They said in the Google Chromium bug list, they said: "We should consider removing built-in support for FTP from Chrome and move it out to an app. Over a seven-day period, only 0.1 to 0.2% of users" - and frankly, I was surprised it was that high - "end up navigating to any FTP URL," and it says, parens, "(with slightly higher numbers

among Linux desktop users). This has been fairly stable over the last year, so it doesn't look like there are trends for FTP to disappear altogether. With the combination of the sockets API and the downloads API, it may be possible to construct a Chrome app which handles this well. Also would need a way to be able to register an app/extension to handle a particular URL scheme so that navigations would be seamless for users of FTP apps."

So they're talking about moving it out of the underlying Chrome browser and into an add-on, an extension, an app for - so that people who did need it could still get it. And again, Leo, I would imagine, especially Linux people, they're going to have nine other ways to grab FTP stuff. And then the little Chrome mention ends, saying: this isn't urgent priority, but might be nice to clean up some code for a little-used feature."

Leo: Yeah.

Steve: So, yeah. So anyway, this sort of just shows the world's changing. It's growing up. And, I mean, it does make sense to shed protocols.

Leo: They took out gopher:// a long time ago. Or did they? Actually, I don't know.

Steve: You know, I think in my entire life I may never have actually used a gopher URL. I knew about it; but, like, eh.

Leo: You can still use Lynx in FTP, by the way. So there. You know Lynx, L-Y-N-X?

Steve: Oh, yeah, the L-Y-N-X.

Leo: The old, you know, the command-line browser?

Steve: Browser, yeah.

Leo: It works great on our site. I was just checking. It works.

Steve: Yeah. And I actually use Wget all the time.

Leo: Wget's great, yeah.

Steve: Yes. You can say, you know, infinite retries, and continue, and aborted retry. And it just sits there as patiently as it needs to be, you know, obtaining a resource, despite problems. Speaking of problems, I was talking last week, I think, or the week before, about - when we were talking about the sale of IPv4 space, and a little bit about routers, the big iron routers, not the little blue boxes that we have on our WAN to LAN interface in our homes, but the big iron routers out in the major telco providers that ship all of the,

like the Tier 1 and Tier 2 and Tier 3 guys, our ISPs and up. And how they have routing tables which, when packets arrive, the packets are inspected from the most significant byte down in order to determine where to send it.

And it helps if geographic regions stay coherent because then it means anything with the same first byte is all going to India or to Europe or wherever. So it makes routing way easier. That means that with only, for example, 250, because they're not all available, with only that many entries for the first byte, you could immediately, if you didn't have to look any further, send stuff on their way. But as we've said, as IP space becomes fragmented, as people with large allocations are tempted to sell off theirs to someone else, somebody who might be in a completely different region of the world, suddenly all the routers need to deal with this.

Well, the way that happens is a protocol we've never discussed in detail, just because it's kind of way at the techie end, and it's known as BGP. That stands for Border Gateway Protocol. BGP is the way routers communicate among themselves to essentially advertise routes that they have to offer and to obtain the routes that their connected routers have to offer. So it's sort of like a big peer-to-peer network. That's very much like what it is. So all the routers are participating in this big, peer-to-peer network, sharing the information about what they know. And that's how the Internet's global routing tables are synchronized. And every so often a mistake is made.

A few days ago, on June 12th, at 08:43 UTC, a mistake was made, a biggie. A telecom in Malaysia, Telekom Malaysia, T-E-L-E-K-O-M, these big telecoms have one or more what's known as "AS" numbers, autonomous systems numbers. And they are AS4788. And this is the way they're known up in the routing world. They, for purely a mistaken reason, started to announce 179,000 routing prefixes for Level 3. Which is to say that they said we essentially route for Level 3 all these prefixes.

So what happened is that suddenly, as this propagated, everybody who had been sending traffic in the proper direction, toward Level 3, started sending it to Telekom Malaysia. That completely buried them. Saturated their connections. Packet losses hit the ceiling, went to 100%. I mean, it basically created an Asian-centric Level 3 outage that was not Level 3's fault. It was because somebody who was a peer in this BGP, this essentially peer-to-peer routing table synchronizing network, made a mistake. They said, yeah, we'll handle 175,000 IP ranges, prefixes, networks for Level 3. And so everyone said, oh, fine, you're closer, here you go, and sent it to them. So anyway, I just...

Leo: It's amazing - this has happened before. It's amazing how easy this is to do.

Steve: Yes. It's a little worrisome. And what's funny is the border gateway protocol, I've been looking into it because I thought it would make a great topic for the podcast? It originated on three napkins. And the designers, literally over a meal, I don't know if pizza was involved, but there were three napkins where the very smart guys in the very beginning sketched out how this would work. And they always intended to replace it. This was an ad hoc, I mean, literally back-of-the-napkin, I assume they flipped them over, sketch for something that they could temporarily do as a stopgap measure, just to kind of have something, to get something going. And as has always happened with the Internet, we're still using it today. And because it was brilliant. They did it right. But they never built it to last. And now we're beginning to see that it is a little creaky. It can have problems. No one generally does this maliciously, I mean, although that's also been a concern. I mean, generally it's a robust system.

But in this posting on bgmon.net, it's on their blog, I'm sorry, bgpmon.net, B-G-P-M-O-N dot net, they said: "This event resulted in significant packet loss and Internet slowdown in all parts of the world. The Level 3 network in particular suffered from severe service degradation between the Asia Pacific region and the rest of their network." And then they posted a graph. They said: "The following graph shows the packet loss - often hitting 100% - as measured by OpenDNS between London over Level 3 and Hong Kong. The same loss patterns were visible from other Level 3 locations globally to, for example Singapore, Hong Kong and Sydney."

So again, an innocent mistake. A major provider said, yeah, we've got a good connection to Level 3. They weren't right, but everyone believed them because, I mean, the routers believed them and sent all their traffic there and completely buried them. And let's see, it began at 08:40, looks like it ended at 10:40, so about two hours. And you can imagine alarm bells were going off. And I'm sure it took someone a while to figure out, what, is this some unbelievable distributed denial of service? That's what it would look like. Suddenly, everybody's traffic is coming to them, and they had to figure out it's because they asked for it. Please, send us your packets. Anything going to Level 3 we'd be happy to have, even though no way do they have the network capacity to actually make good on the claim that their router was broadcasting.

Leo: Wow.

Steve: Just amazingly, amazingly cool. And Wikipedia has gone all HTTPS. They announced, I think it was on Friday, that they were making the switch. For a long time you could use HTTPS Everywhere, that is, you could switch to HTTPS; but they weren't switching you, us, users of Wikipedia to HTTPS. The Wikimedia Foundation blog posted: "To be truly free, access to knowledge must be secure and uncensored. At the Wikimedia Foundation, we believe that you should be able to use Wikipedia and the Wikimedia sites without sacrificing privacy or safety.

"Today we are happy to announce that we are in the process of implementing HTTPS to encrypt all Wikimedia traffic. We will also use HTTP Strict Transport Security, HSTS, to protect against efforts to break HTTPS and intercept traffic. With this change, the nearly half a billion people who rely on Wikipedia and its sister projects every month will be able to share in the world's knowledge more securely.

"HTTPS," they wrote, "is not new to Wikimedia sites. Since 2011, we're working on establishing the infrastructure and technical requirements and understanding the policy and community implications of HTTPS for all Wikimedia traffic, with the ultimate goal of making it available to all users. In fact, for the past four years Wikimedia users could access our sites with HTTPS manually." Or, as I said, through HTTPS Everywhere. And so they said: "Over the last few years," they conclude, "increasing concerns about government surveillance prompted members of the Wikimedia community to push for more broad protection through HTTPS. We agreed, and made this transition a priority for our policy and engineering teams."

And so this morning I went to Wikipedia and grabbed the domain name and went over to SSL Labs. They get a very strong A with the way they have, at least the server I'm seeing, which was en, for English, .wikipedia.org. They are using only the top TLS protocols, 1.0, 1.1, and 1.2. They broadly support Perfect Forward Secrecy, which is to say that all of their first ciphers in their cipher suite are ephemeral key Diffie-Hellman, which is exactly what you want. They've got the key strength backwards. They have 128-bit strength in front of 256. I don't know why. But I would have swapped those, and do,

on my server. But still very good.

And overall, very nice cipher suite ordering, again, on the server that I saw. And they're offering OCSP stapling, which is the way you want to do it. You want the server to provide the OCSP, that is, the certificate revocation information, in the handshake. That way browsers can then, when they see that the server is stapling, they can do a hard fail, and that's the way you're able to get the benefit of state-of-the-art revocation, I mean, robust, no-way-to-hack-it revocation, at the same time as not inconveniencing anybody with an OCSP server that is not responding.

And the one thing they did, and I let them off the hook for this, but SSL Labs noted that their HSTS duration was a little short. It was - I've got mine set to maximum. I don't even know how many years it is. Theirs is set for 180 days, so half a year. So what that means is that, if you didn't visit within half a year, then your browser's knowledge that Wikipedia is only secure would expire. And so there's a tiny little window of vulnerability, nothing to worry about. And maybe they just didn't set it to maximum, I'm not sure why, because 180 days, it sort of might as well be forever in terms of usability. But they chose a shorter time, not maximum. And SSL Labs did note that. So, yay. We have secure, you know, enforced security on Wikipedia and on the Wikimedia properties. So again, just more of this movement in this direction.

So a couple miscellaneous things. I need help. The Screen Savers, The New Screen Savers show has asked me to provide, to pre-record some tips of things that would be dropped into the show when it makes sense, and just to provide some variety. And it's a classic case of being too close to the forest, or being a tree, I guess, because I can't think of a thing. I did one about the problem with car fobs and keyless entry.

Leo: That's a good one, good.

Steve: So you have that.

Leo: Thank you.

Steve: But I'm dry. And then I thought, you know, I know who to ask. I should just ask our listeners. I'm sure you guys...

Leo: What would you like to know?

Steve: From your perspective, or what do you already know that a different audience should know? And I just can't come up with anything. So tweet it to me or send it through GRC.com/feedback, and put something like TNSS in the subject, The New Screen Savers - in the same way that this is SN, that show is TNSS - because they're snippets for things that I could prerecord, like about a minute long, not very long, that I could help people with because I can't think of anything to say. I mean, I guess the problem is I've already said it. I'm happy to say it again, but I just sort of need some ideas from our audience. So this is an open call for anyone to send me their ideas of what they think would be a great little short segment to drop into Leo's newest and hit show.

Leo: Good. Good.

Steve: I already mentioned that the new Syfy show, "Dark Matter," that's the name of it, looks awful. A crew of bad actors wake up with amnesia in deep space and don't know why they're there and don't know who gave them the job because they can't act. We do have a show starting toward the end of this month called "Humans" on AMC that looks interesting. I just thought I'd mention it. It's not about humans. It seems to be about androids. But I don't know anything, except I just wanted to make sure people knew about it. And then I did want to mention sort of a weird thing Twitter's decided to do. They've decided to allow direct messages to increase their length to 10,000 characters.

Leo: Yeah, I don't like that at all.

Steve: It's, wow. Yeah.

Leo: They're trying to turn into an email client or something.

Steve: I would like more than maybe 140. I mean, it's a little bit annoying to have that limit. But also it's sort of a blessing to have that limit. I mean, I wonder if Twitter needs to change at all, if it isn't just fine the way it is, even though they can't seem to make any money, and I hope they somehow stay alive. One thing that really annoys me is I keep telling them to stop sending me email, and they ignore my settings every so often. They just turn them all on again and starting sending me crap. So I have to go back and deliberately turn it off again. I guess that's the price I pay for it being free and for it being a service that I really do depend upon. But anyway, so for what it's worth, I'm not looking forward to 10,000-character direct messages.

For my SpinRite mention today, I don't have a testimonial. What I do have to share is some interesting research which was published in a very detailed, fact-filled, 14-page Carnegie Mellon research paper with Facebook, latest research on SSDs. And they said, just the abstract from the paper, they said: "Servers use flash memory-based solid-state drives (SSDs) as a high-performance alternative to hard disk drives to store persistent data. Unfortunately, recent increases in flash density have also brought about decreases in chip-level reliability. In a data center environment, flash-based SSD failures can lead to downtime and, in the worst case, data loss. As a result, it is important to understand flash memory reliability characteristics over flash lifetime in a realistic production data center environment running modern applications and system software.

"This paper presents the first large-scale study of flash-based SSD reliability in the field. We analyze data collected across a majority of flash-based solid-state drives at Facebook data centers over nearly four years and many millions of operational hours in order to understand failure properties and trends of flash-based SSDs. Our study considers a variety of SSD characteristics, including the amount of data written to and read from flash chips; how data is mapped within the SSD address space; the amount of data copied, erased, and discarded by the flash controller; and flash board temperature and power consumption. Anyway, a really interesting paper. And what they found in summary is some good news and some bad news. About temperature, they discovered something that we have just been talking about recently about SSDs and temperature. So they confirm that SSDs are sensitive to temperature, more so than hard drives that

really don't care. When they get hot, the SSD may throttle back its performance. And so they rhetorically asked: "Experiencing unexplained slowdowns on some servers? Check the temperature. First-generation SSDs failed more often as temperature rose, possibly due to a lack of throttling. Some second-generation SSDs throttle aggressively enough to reduce their failure rates, while others keep the failure curve flat." So again, temperature is crucial.

And they found that SSDs are power thirsty. The PCIe v2 SSDs ran anywhere from 8 to 14.5 watts of power, which they said they felt was a high and surprisingly wide range. The team found that, as power consumption rose, so did failure rates. And then of course write fatigue, which is the concern with this technology of nonvolatile RAM storage. They found that the level of system write activity correlated directly with SSD failure, probably because flash writes require a lot of power. That is true because, remember, they're essentially - they increase the power in order to essentially break through insulation in order to deposit or remove charge that are stranded on the other side of an insulator. Because, basically, they're just an array of capacitors. So they said disks could be a better choice in heavy write applications such as logging, rather than SSDs.

And then about failures, they described SSD failures as so-called UREs, Unrecoverable Read Errors. They said in SSDs they are relatively common. Between 4.2 and 34.1% of the SSDs reported uncorrectable errors, so as many as a third. In fact, 99.8% of the SSDs reporting an error in one week reported another error in the next week. So once they start to have problems, they typically have more problems.

And so they said, of life cycles and failures, the SSD failure profile differs from disks. Where disks exhibit an infant mortality effect, then they enjoy a few years of good reliability before becoming less reliable. SSDs have an early period of unrecoverable errors as faulty cells are identified, and then their reliability increases until cell wear-out leads to increasing read failures.

And then the last thing they found was that the way data is written to them affects it, that is, the data layout. They found that SSDs do not like a sparse data layout, and they're much happier with large block, large contiguous activity. And of course we know why. And that's because the SSDs themselves have an internal blocking architecture; and, if you change just one bit in one of those blocks, it's necessary for the SSD to read the entire block, write it all to one state, and then rewrite it back to the given state, the original state, with that one bit changed. So the technology difference means that you really, if you can, you want to write as much within a single contiguous region as possible. Otherwise, you're actually wearing them out in a way that disk drives have no similar problem. So really interesting.

And of course all of this, and the prior experience of our own listeners who have reported that SpinRite has been effective for them in recovering SSDs, bringing them back to life from dead, just as it does with hard disks, is one of the reasons that I have such a good feeling about SpinRite's future and intend, after doing 6.1 and 6.2 and 6.3, that v6 series which will be free for everyone, I plan to then, the next thing I do, probably, subject to other things that do tend to arise like SQRL did, move on to a complete rewrite and do a v7. I think SpinRite has a long life ahead of it.

Leo: Congratulations. It's nice to know.

Steve: Yeah.

Leo: It's also nice to know SSDs are at least a reasonable alternative to spinning drives. I mean, spinning drives have many of these same problems, obviously. Different kinds.

Steve: If nothing, I would say they're fast. They read fast. That's probably their advantage. See, the problem is, we've talked about it before, I managed to get on my servers single-level memory, you know, it's called SLC, Single Level Cell, versus MLC. You can't get them anymore because Multi Level Cell, where you store multiple voltages in that little capacitor, that allows them to get two bits, or three bits, or maybe four bits in the same capacitor by storing different voltages. The problem is then any variation in that reads back as an error. So you're squeezing the tolerance down. And of course the density's been going up, both by doing multilevel bit storage and making the bits smaller. If you have a smaller capacitor, then the charge has fewer electrons on it, so there are fewer to be lost in order for it to change its value.

So, I mean, the problem is, exactly as this paper said, unfortunately, the push for density, to bring costs down and density up, has pushed SSDs out into the same performance reliability envelope as hard drives. So, yes, they read faster. But once upon a time we had this idea that, oh, they're solid-state, so they're going to be more reliable. Turns out all of our experience and this paper demonstrates no. They die a little differently. Different things make them, you know, like, hurt them; and they actually, you know, hard drives can wear out just from age. But the actual storage medium of SSDs wears out with age. So anyway, I'm glad I will be addressing them directly with SpinRite.

Leo: We continue with Security Now!. Steve Gibson, let's talk tracking.

Steve: Okay. So in Firefox v35, Mozilla added a sort of an experimental feature, which we talked about at the time, called Tracking Protection. And they did not surface it at any user interface except for the about:config, the treasure trove of Firefox's configuration widgetry. So it's off by default. But users who are interested to experiment with it could turn it on. And we are told that a future version of Firefox will bring a simple checkbox out to one of the tabs on Firefox's config panel in the UI, making it something easy to turn on and off.

A gal who has since left Mozilla, she was there until the beginning of April, Monica Chew, was responsible for examining the whole privacy side of Firefox. And she and a coworker of hers wrote a paper titled "Tracking Protection for Firefox," and presented it at the Web 2.0 Security and Privacy 2015 Conference. She said - I'll just quote from the top of this paper. "This paper is the last artifact of my work at Mozilla, since I left employment there at the beginning of April. I believe that Mozilla can make progress in privacy, but leadership needs to recognize that current advertising practices that enable 'free,'" she has in quotes, "content are in direct conflict with security, privacy, stability, and performance concerns, and that Firefox is first and foremost a user-agent, not an industry-agent.

"Advertising does not make" - and I liked what she said here. She said: "Advertising does not make content free. It merely externalizes the costs in a way that incentivizes malicious or incompetent players to build things like Superfish, infect one in 20 machines with ad-injection malware, and create sites that require unsafe plugins, take twice as many resources to load - all quite expensive in terms of bandwidth, power, and stability."

And then she concludes, saying: "It will take a major force to disrupt this ecosystem and motivate alternative revenue models. I hope that Mozilla can be that force."

So I think she's clearly a privacy and security evangelist, maybe overstating a little bit, like bringing Superfish in, for example, as regards to tracking. But it was interesting that she has this stat that one in 20 machines are being infected with malicious, ad-injected malware. But I wanted to share with you, Leo, to sort of discuss where this stands relative to, I guess, Firefox versus Google. One of the things that I thought was so interesting about your discussions that you've had in the last week has been about how Apple differs from Google in terms of Apple not selling advertising, but Apple selling hardware. And so their privacy policies are different than Google's because, essentially, people are sort of paying upfront for the experiences they have with Apple machines, where Google is saying, no, let's sort of have a pay-as-you-go through an advertising-based model. And I think Leo has left the room.

Leo: No, I'm here.

Steve: Oh, oh, sorry. Yeah, so, and I guess my sense is that Firefox and Chrome are sort of in the same sort of position, where Firefox is saying, you know, we're going to experiment with blocking tracking because we're going to take the position that the browser is the user's agent, not the industry's agent.

Leo: I'm sorry. Do you need me for something? I'm sitting here, but I don't hear a question in all that.

Steve: Oh, okay. Well, okay. So...

Leo: What would you like to know, Steve?

Steve: I just thought you'd have an opinion.

Leo: How I feel about that?

Steve: Okay. So...

Leo: I don't know. I mean, you heard the debate we had on Wednesday, obviously.

Steve: Right. So I guess my feeling is malware is a problem. Scripting is a problem. And, like you, I'm not that concerned, I'm just not that concerned about tracking and profiling. Personally, it doesn't worry me that much. But at the same time, I fully respect the rights and sentiments of those who are concerned over the practice. And many of the guests that you interact with, Leo, on your podcasts, demonstrate a great deal of concern. We know that our listeners do, and people in the chatroom do. So there's a population of people that do. What bothers me is intrusion into my browsing experience. That is, cover-up ads that require me to take action. I was watching you a couple weeks ago, you

were looking at, I think it was the USA Today site, and they were, like, running videos, like had videos playing on the pages.

Leo: So obnoxious, yeah.

Steve: But I can't remember what you were looking for. But, I mean, it was just like...

Leo: It was an autoplay ad. And you're seeing that more and more. But, you know, that all comes from people blocking ads; right? I mean, it's kind of a vicious circle because you block ads. And then the advertisers say, oh, the ads aren't working. So now let's make it more egregious. Then you say, that sucks. I don't want that. So you block more ads. I don't know. That's, you know, that's not a good...

Steve: Yeah. The thing that made me turn scripts off on Safari, on my iPad, was when I went somewhere, and the screen was covered up, and there was a notice saying "Change your device orientation to proceed." So...

Leo: Change your device. So you have to turn it sideways?

Steve: So, yes. This ad was not happy with the way I was holding my iPad. And I said, okay, no, no, I just, you know, no more. So I turned - and there's been some inconvenience associated with it. But I'm willing to do that. So, oh, and I...

Leo: It's actually good for us. I have to point out, the reason that we do well is because we don't do that kind of thing to you. And the ads you hear on TWiT are interruptions, and you can always skip them, of course. But they're not - we don't turn you upside down and shake you.

Steve: And you don't have me holding a Harry's razor and shaving on-camera.

Leo: But in defense of ad-supported free media, you know, the option is to put a paywall up and say you have to pay for this content. And I don't really want to do that.

Steve: No. And in fact, I have in my notes here that paywalls don't work for me because, for example, I'm not at The Financial Times or The New York Times...

Leo: Right, you just skip it.

Steve: ...often enough to subscribe.

Leo: Right.

Steve: And so when people send me links that go to FT.com, it's like, I'll go there. It'll say, oh, here, you know, happy to have you subscribe. It's like, no, I just - can't I see this? So, for example, the idea of a limited number of pages per month, that really does work for me.

Okay. So here's the problem, though. I believe that the technology, the nature of our model is broken. And can you switch so I can see you? Because I feel I'm wanting to talk to you about this. Thank you.

Leo: Okay, sorry.

Steve: Yeah, I just - because I feel like we're in a mode where we have, like, perverse incentives. For example, a site doesn't host the content itself. There's no bandwidth cost for it. All it does is have links that require our browser to go get something else. And it can't even represent with any authority what the browser's going to get because it has no control over that. So, for example, it might be a link to malware. And this is how advertisements, or malvertising, how advertisements are infecting people, or Flash is infecting people. I mean, it's Flash on sites that are running ads that are causing problems. And so the site itself is just putting links on their page that cause the people who visit the sites, browsers, to go get these things. And as you said, there is this tendency to escalate, also.

Leo: Yeah, that's called - that's a link farm; right? You're talking about a link farm, basically.

Steve: Well, actually, I mean, that's what sites have become.

Leo: Well, people put some content on the site; right?

Steve: No, well, not a link farm, but a reference farm. They're putting, like, script tags and image tags and all these things to third parties, so that they're not serving the content.

Leo: When you're talking about ad - I'm confused. Give me an example. Like USA Today?

Steve: Exactly, like USA Today.

Leo: Okay. So they have their content, which they put on there. But there's also ad content, and that isn't, as you say, and our ads aren't, for good reason, served by those sites. They're served by ad-serving sites.

Steve: Well, not for good reason. It's for convenience.

Leo: Oh, yeah, trust me for good - no. It's for good reason. An advertiser wants to know how many times that ad was viewed. They're not going to trust our information about that. They want to know what the impressions are, and that comes from a third party. Not us. Otherwise they'd have to just say, "Well, how many people saw your site?" "Oh, a billion. Billion 215.

Steve: Okay. So...

Leo: So that's one good reason. There's also a lot of complexity in the way those ads are rotated. We have, for instance, our ad banners, you'll see a different banner. The number of times you see any given banner is handled by the ad-serving software. So after you've seen a banner five times it says, okay, you got your five impressions. It's very complicated. It's not something we do because - Google does it for us, by the way. That's why you see, if you go to our site with Ghostery or something, you'll see, I think, three different Google tags. Our own Google Analytics, and you'll see probably New Relic, which is a performance metric monitoring service. But then you'll see three Google-related tags. Those are those two ad banners - there's only two ad banners on our whole site. But we can't serve them. That wouldn't work. I mean, I guess we could. We could write an ad-serving software to do it. It's not practicable.

Steve: So, okay. So the problem, though, is that - and no one is having any complaints about TWiT. But, for example...

Leo: Actually, they do, because every time somebody runs - anybody who runs Ghostery or NoScript or any of these things...

Steve: Sees something coming from somewhere else.

Leo: ...sees this stuff, and they complain. But guess what. You're watching ad-supported free media. I mean, there's a responsibility for the companies that do ads not to be bad. And there's certainly a disincentive for them to be obnoxious.

Steve: Well, and so I guess that's - so that's one of the issues that affects us is that there is a lot of content which is coming from third-party sites when we visit first-party sites.

Leo: Right, right. That's very common.

Steve: Yeah. So, for example, what Mozilla's tracking protection does, this thing that's built into Firefox, it's a curate - it's not an ad blocker, I should say. I mean, I've had it turned on for months. And I see ads all the time. And, I mean, and I do feel, as we've said, sort of an ethical responsibility. I mean, I'm not clicking on them. They're not doing

anything. But in order to support the content of the web, I'm happy to have them delivered to my browser. But there are a class of non-ad trackers that just do tracking. I mean, for example, their business model is to generate tracking information, as I understand it. And, for example, it's scripting as opposed to images.

So what the Firefox tracking protection does is it blocks Firefox's fetching of about 1,500 domains, by domain. It's not blocking third-party cookies. It's just saying we want to block this domain. And in the process, in the statistics that they have gathered, it generates a dramatic improvement in performance. For example, in the paper where they run the stats, you go to Weather.com, and the page loads in 3.5 seconds if you block the additional trackers that they load, or 6.3 seconds without. And it pulls 2.8MB if you have tracking blocked, versus 4.3 without, and more than double the number of HTTP requests. So anyway, I guess the only thing I really had to add to the discussions you guys have been having is that it feels like there's this tension in the industry between the advertisers and the users. And I guess it's the fact that few enough people use ad blockers that it just doesn't matter that much.

Leo: Yeah. We can pretty much ignore it at this point. However, if it got pretty widespread, might make a difference. You're seeing this in broadcast television already. They're doing more live - look, you're getting crap on broadcast television. There's a reason. The reason you bemoan this lack of good sci-fi, because people skip ads. So they do live stuff, "American Idol," because then people are incented to watch live so they can vote, and they can't skip ads. You get what you pay for, folks. And you can skip ads, but people just will abandon you. If they can't afford the production, they're not going to do it.

Steve: In their look at the industry, they had an interesting stat. In 2013, ad blocking grew nearly 70%, with nearly 41% of the 18 to 29 year olds reporting usage. So among...

Leo: So what you're going to see is more sneaky ads. You're going to see more ad placement, product placement. You're going to basically push them underground.

Steve: Yeah. In fact, I saw...

Leo: So I'm not against it. If you want to do it, do it. That's fine. But just understand.

Steve: I saw you comment that there was a type of ad, you had a name for it...

Leo: It's called "narrative content."

Steve: Yes, where it looked like, sort of like little windows at the bottom of the page that looked like links to other stuff.

Leo: Oh, that's chumming, yeah. Ad chumming, yeah. That's really repulsive. I mean, I understand, I mean, we brought it on ourselves because we - not we, me, but content creators brought it on themselves, and advertisers brought it on themselves, by going overboard.

Steve: Well, I mean, and it is a means of monetizing. I think one of the tricky things is it's a bit of a slippery slope because someone who has a website with ads, I mean, it takes a great deal of self-control not to add another one because you look at the amount of revenue that your existing ads have and say, gee, you know, if we added another one, we'd make more money. And so, I mean, to me, it's interesting from a technology standpoint. And as a user, it'll be interesting to see how this evolves over time because it feels to me like what we're seeing is sites, other than yours and mine, are apparently offering scripts which generate revenue from services that compile dossiers, not displaying ads, but are in the business of tracking users. And so they're scripts that produce no visible content, have no value to the user, and are only being used to track them.

Leo: Well, who would do that unless it was to give them better ads? There's no reason, otherwise. So ultimately it's an ad service. There's no - why would - yeah. If that's the case, it makes no sense. I don't understand why they're tracking them, what they want. There's two reasons why you track somebody, to give them better, to give them ads that are more relevant to their interests, or if you - and we track users, yeah, because when you come in, you get Google Analytics.

Steve: Well, exactly. You want to profile who your visitors are.

Leo: Well, no, we don't profile them. We just want to know how many there are. There's no profiling going on in our...

Steve: Well, but you don't.

Leo: But how do you distinguish? You don't know what the code is getting, do you?

FEMALE VOICE: Showing pictures matching your request.

Leo: Sorry, I said the word "Google."

Steve: No. So sites could certainly be - sites could buy profiling services.

Leo: I see. And sell that information.

Steve: And they put code on - exactly. So they put code on their site to allow a third party to track, to aggregate information about all the other sites that their site's visitors

go to, and pull this together and say, here's the demographics of the people who are looking at this page and this page and that page.

Leo: The only way that's of value is if you're an advertiser.

Steve: Or if you're the person running the site.

Leo: Or the site, yeah, the site might want to know.

Steve: And you want to know the profile of the people who are visiting.

Leo: But, I mean, ultimately this is about advertising. I don't know who else this would be about.

Steve: Or, well, or someone who wants to collect intelligence on the demographics of a site.

Leo: But what's the purpose of that? For advertising.

Steve: Optimizing the content.

Leo: Yeah. But in both cases, I don't understand why that bothers people so much. But so how do I - you can see I haven't used Firefox in a while. My home page is Protopage.

Steve: Ah, good. So go to about:config.

Leo: By the way, you can bookmark these, somebody in the chatroom is mentioning.

Steve: Yup.

Leo: Config, okay.

Steve: And hit Enter.

Leo: Okay.

Steve: And you'll get a bazillion things.

Leo: This might void your warranty? I'll be careful. Let me uncheck that box. Thank you.

Steve: Interesting. Okay. Now, in the search box, put "tracking."

Leo: And as somebody pointed out, you can actually get a URL, I guess, for this. And so which one should I...

Steve: And there ought to be "tracking protection."

Leo: Yeah, there are several.

Steve: And, like, at the third one down, that one.

Leo: All three of these. Okay. It's enabled, right.

Steve: Okay. That one, double-click to turn it on.

Leo: Oh, yeah, we want it to be true. Okay.

Steve: You want it to be true. Tracking protection enabled. And now go look somewhere. Go anywhere you want.

Leo: Well, first let's go to our site. What would it do on our site?

Steve: No effect. That's what I'm saying, is this is not an ad blocker. It's blocking scripts that you're not seeing.

Leo: USAToday.com. Since they're the poster child.

Steve: Yup.

Leo: Oh, I don't see a lot of crap.

Steve: Correct.

Leo: So it's blocked invisible stuff, in other words.

Steve: Correct. I went to USA Today this morning. And without blocking, my browser made 162 requests and pulled 4.5MB of data.

Leo: Whoa.

Steve: I turned blocking on, and it made 139 requests and only 3.6MB. So that meant that there were 23 tracking domains and another meg of data.

Leo: That is a lot.

Steve: And these are typically big JavaScript libraries because, again, these tracking - they don't care. They have a CDN which is providing this. And my browser is bringing down a meg of JavaScript in order to be tracked by something that's not enhancing the content. I mean, it's not visible. There's no visible difference on the site because these are just tracking domains. And in their statistics, they looked at the top 200 news-oriented sites, and they said, "We demonstrate a 67.5% reduction," so nearly two or more than two thirds in the number of cookies set during a crawl of the Alexa top 200 news sites, and a 44% improvement in performance. So again, as you said, here was an extra megabyte of stuff downloaded with tracking enabled versus not. And as you see, now you've got...

Leo: That's a lot, yeah. The better choice would be just not to ever visit that site. I guess there's no way of knowing ahead of time.

Steve: It's all sites. I mean, it's all over the place.

Leo: You vote with your feet.

Steve: Right.

Leo: Well, yeah, look at that. I turned it back off, and now I'm getting a big ad that popped up. That was that thing. Now I have to press a button that says "Get the news" to get rid of that ad. Oh, there's the site I saw without the tracker, without the blocking. Hmm. Well, kind of my attitude is...

Steve: I do, I mean, I understand.

Leo: You probably, best thing to do would be never visit those sites. Like as soon as you realize there's a site that's drawing megabytes of data behind the scenes, just stop using USA Today.

Steve: Right.

Leo: Right?

Steve: Right. I mean, I think we have - we have, I mean, where does it end? How does this - where does it go, is what I'm curious to see, because this has been sort of incremental. It feels like the load of ads is growing on sites. And I would argue, Leo, that it's more just greed from sites than it is people blocking ads. I think the percentage of people who block ads is minimal. I mean, just vanishingly small. I don't know anybody else who does it except some percentage of our listeners, probably.

Leo: Yeah.

Steve: And other techies. And I've heard you say on other podcasts that we know that the effect of ad blocking is negligible. So the problem is, if you have a bunch of ads on a page that is generating money, and you have a slow month, it's hard, it takes a great deal of self control not to put some more ads on and have more revenue as a result. And at some point the site becomes unusable. I don't know.

Leo: So this is disabled by default.

Steve: Correct. And no good UI, no user-convenient UI at the moment. It's supposed to be coming soon, so that you would be able to just go into the standard config under Privacy and turn on "Enable tracking protection."

Leo: Yeah. I don't have a problem with people turning that on. I'm sure USA Today and Huffington Post do. But I don't have a problem with it. I don't know what the moral implications are. It's kind of like saying - there is some ethical consequence to this.

Steve: Well, hey, what about - I meant to say on the podcast, when you guys were talking about it, what about the Reader view? Like in Safari, it's got that little Reader view. You click it, and it cleans up the page.

Leo: It's nice.

Steve: Which basically just gives you, well, it's nice, but what does it do?

Leo: Same issue.

Steve: It just gives you the content.

Leo: Right, same issue.

Steve: Yeah.

Leo: It's kind of like saying, you know, this only comes up because we can do it. We don't go into a McDonald's and say, you know, I really hate the color scheme here.

Steve: Good point. It's using technology, and we have control over the technology.

Leo: Right.

Steve: One of the other interesting things is there was a quote from Google, and I'm not seeing it in my notes, but Google said that a huge percentage of ads are never seen.

Leo: Right.

Steve: And of course that's because they're so-called "below the fold."

Leo: Only a third of ads are seen for a second or more.

Steve: Right, because if you don't scroll down, you're not going to get them.

Leo: Yeah, we can't charge for the ads below the fold. We don't charge for that banner at the bottom of the page.

Steve: So it knows if you scroll down and see it?

Leo: No, we just don't charge for it. It's because we figure people won't. We give it as a spiff to advertisers.

Steve: Oh, nice.

Leo: The only ad we charge for is that little thin one at the top there. You know, what can I say. You and I are sitting here. You don't live on ads. I do. You take away the ads, I'll have to get a job. And I'm not the only one. And in fact the state of journalism in this country is going downhill very rapidly. And I would submit it's because, I mean, and I'm not saying that you're to blame, or anybody's to blame for that, because obviously the ad industry is really obnoxious.

Steve: I just don't think the model works. I mean, even if everyone - if there were no ad blockers, I don't click on ads on web pages. I'm not - they're not getting any revenue from me. And a lot of them are, as you said, are obnoxious. I mean, I just - it feels to me like there's a problem.

Leo: The ethical thing to do would be not go to sites like that. I understand you don't know ahead of time. So the right thing to do would really be for Firefox to put up a little thing that says you just, I mean, I'd love to see that. And somebody should write a plug-in that says "This site costs you an extra 5MB and 30 extra seconds with invisible trackers."

Steve: Ah, "Would you like to blacklist it?"

Leo: "Would you like to blacklist it."

Steve: And that way you can say, yes. Remind me, next time I click on a link, give me a little popup and say, oh, you said you don't want to go here. It's like, well, oh, but I really wanted to look at this.

Leo: Well, that's the ethical conundrum because you're using content that's paid for by those things.

Steve: Yes.

Leo: And what you'd like to do is to use the content and not pay for it. And so I have an ethical issue with that. I mean, I think the right, ethically, the right thing to do would be just not to go to that site.

Steve: Well, okay, wait a minute. I'm happy to look, I'm happy to have the ads, as long as they don't make me rotate my device to a different orientation.

Leo: I agree. I agree. So you shouldn't go back there. You don't get to vote, really, on how a site does its business. You get to vote with your feet, but you don't get to modify the site. I mean, in fact you can. It makes people do it. But it seems to me unethical.

Steve: Except that no one is going to build a list of sites in their head that they know not to go to.

Leo: Right. So if I...

Steve: You're always going to...

Leo: So I would like to see Firefox modify this so that it just lets you know and gives you an option to say don't go there anymore. That would be appropriate. And by the way, the ad industry would hate that, too. But I think that that's a more ethical point

of view. What I consider unethical is saying, fine, I'm going to technologically remove all the revenue points on your site and still consume your content. That's stealing. Right?

Steve: Yeah.

Leo: I mean, not that I don't do it. I skip ads. But by the way, notice there's very little good content on ad-supported network television. It's mostly live stuff.

Steve: Yeah, you and I are DVR users, and I do not watch live television. I wait for it to go onto the drive, and then...

Leo: Guess what you're going to get? You're going to get only live television. You're going to get basketball games and "American Idol," and you're not going to see produced stuff except on HBO, where you have to pay for it, and places like that. Because if you break the ad model, that's what happens. Now, you say the ad model's broken, and maybe it is. That may very well be true. I mean, there's a lot of evidence that the banner ads don't work for nothing.

Steve: Yeah.

Leo: So I just think it's unethical to say, I don't like how you're doing business, so I'm going to - it's like stealing a candy bar. I'm going to consume your content without paying for it. That's stealing.

Steve: Yeah. But you would not - you have no problem with anonymizing, that is, with not being tracked. Or do you think we need to be tracked also?

Leo: If that's how a business sets it up.

Steve: Apparently there's value in tracking.

Leo: There's value in it, so that's how - and, by the way, you know who makes money on the web? About five sites, like Huffington Post and USA Today. Most of the rest...

Steve: Well, and Google.

Leo: And Google. And the rest of them are going behind paywalls, which doesn't work. So this is - you're going to have - there's a consequence to this behavior.

Steve: Well, but the behavior is not clicking on ads; right?

Leo: Well, I don't click on ads. Of course. I don't answer spam. But this isn't spam. This is content you want to read. I'm not saying you should. You know what I'm saying?

Steve: The non-advertising content is content we want to read.

Leo: People, I mean, look, there's a lot of things you could say. We think everything's free because the web seems to be free. So we don't really think about the fact that this isn't free, far from it. I spent a quarter of a million dollars on our new website out of my wallet right here. It came out of this wallet right here. So...

Steve: And it's not that I'm not ad supported. I make a point of telling people about SpinRite every week.

Leo: Right. I have to monetize. And we send you a check.

Steve: And it's the only reason I'm here. It's the only reason I'm here is that works.

Leo: You could return the checks, and we'll take the ads off your show. I mean, I just - it's like - it's a challenge to me. But I just really want to raise that issue of the fact that, if you don't like the way a site does business, don't visit the site. What you shouldn't...

Steve: See, that doesn't work either because you've already visited.

Leo: Well, that's why I would love to see a technological solution that lets you say, oh, good, fine, I don't like all this, block the site. But what really people want to do is steal the candy bar. They want to go to the site...

Steve: That's correct.

Leo: ...and block all the revenue. Well, I don't think that's ethical. Doesn't affect me, by the way. People think, oh, Leo's saying that because he's ad supported. It doesn't affect me. What we do, and one of the reasons we do it, you know, do you want to know why there's not, why TechTV doesn't exist? This is why. Technological people do this stuff. And nobody in their right mind makes content for technologically sophisticated people. And we are very - we have to be very, very, very careful because I know I will hear from every one of you, rightly so, if we put a tracker on there. We'd hear from you. I hear all the time about, oh, my Ghostery report says - I hear that all the time. So by the way, that's why there's no TechTV. Very simple.

Steve: Yeah.

Leo: So now think about it. Now what do you want to do? I'm not saying there's an easy answer. There isn't.

Steve: Yeah, I don't think there is. I think the answer is to, I mean, I guess my question, or my issue is the idea of running scripts on a site because I don't want my computer infected with malware.

Leo: I completely understand. Or waste bandwidth supporting something you don't even know about. I understand that.

Steve: Yeah, that is hidden JavaScript used for building a profile of me.

Leo: Right.

Steve: And certainly I'm not obligated to accept malware from a site that I visit that has an infected...

Leo: I completely agree.

Steve: So, I mean, but people do get infected, one in 20 apparently, from malicious advertising.

Leo: Look, I didn't want to have any banner ads on the new website. I was told very clearly that was not an option.

Steve: Yeah.

Leo: I was told very clearly that's not an option. What can I say?

Steve: Yeah.

Leo: I don't, I don't, I don't, you know, we have to monetize. We can't - this costs millions of dollars a year to do this.

Steve: I've had, it actually was Mark Thompson, I had Mark Thompson tell me that he's beginning to think that my site looks dumb, not just because of its design, because we have that, but because there's no ads. Like ads are now sort of...

Leo: Well, I don't know if I agree with that.

Steve: Well, he did tell me that, that they're part of being in business, being on the web, is having some non-intrusive ads. It's like, well...

Leo: I could argue that your site, your whole site is an ad for SpinRite.

Steve: Yeah.

Leo: Everything you do that raises goodwill is a very sophisticated way of promoting SpinRite. It's sophisticated. And we had to do the same thing. We can't exist, we couldn't do a Huffington Post or a USA Today-style website because our audience would reject it. So we can't push that envelope too far. But we have to have ads of some kind.

Steve: Yeah, yeah.

Leo: So, I mean, I think that, really, that this is the best, TWiT would be an example of the best-case scenario.

Steve: I think it's a beautiful model, Leo, I do.

Leo: Yeah, where we have to think about what the audience wants. And we really endeavor, you know, we think about this stuff all the time.

Steve: And I get email from Lisa when the network is considering a new sponsor. It's like, hey, Steve, are these people that you would feel good about having as a sponsor?

Leo: Oh, yeah, we always do that, of course.

Steve: And I check them out and say yea or nay.

Leo: The other thing we do, I mean, for instance, and I'm kind of in the middle of this, having just gone through six months of web design hell, which ended nicely, I think. But one of the things that was an issue is the images on our site. Lisa said, "They don't look good when I blow them up." And I said, well, yeah, because we're very cognizant of the amount of bandwidth. And so we want those images to be small. They were about 60 to 80K. And she said, well - so we turned the compression up to, instead of 65% JPG compression, to 95% JPG compression, which balloons the size of these files up quite a bit. And I've already seen somebody say, "On my very slow bandwidth, those images take too long to load." The

question, and what happens all the time, is it's a moving target.

Steve: Yes.

Leo: It's like, well, how much, how big should a website be? How big should images be? Obviously they shouldn't be 5MB. Can they be 100K? Somebody complained because they're 90K. So I don't know. We just try to kind of always listen to our audience. Maybe that's what's wrong. These sites don't really listen. Right? Actually, USA doesn't have to. USA Today doesn't have to. Nobody's complaining to them. Five geeks. That's all.

Steve: Right, right.

Leo: Although I think those huge takeovers, how could anybody like that? I know I don't go to USA Today ever. I very rarely go there for that reason.

Steve: Yeah, and it's funny because, I mean, I didn't even see that because I also have NoScript, just because I don't want scripts. I'm worried about scripts from a security standpoint. So I saw that page that you saw with tracking protection.

Leo: Right. I'll show you the tweet I got, the four or five tweets I got from a guy who said, "Your site looks terrible. Oh, my god, I can hardly read it. Oh, it's the worst site ever. It's terrible." And then he says, "Oh, never mind, NoScript was blocking it."

Steve: Well, yeah. And in fact, I told you that at the beginning of the show, is that I had new.twit.tv. And I guess maybe I've always had script, I mean, I run with scripts blocked by default. And so all of the pictures you have went down, scrolled down the page. And then instead of the little moving...

Leo: Yeah, because you had NoScript on.

Steve: Right.

Leo: Anyway, thank you, bothyhead, for the report. I appreciate that. And even better for saying, whoops, I had NoScript on. Yeah, I recommend, if you want to use most websites except Steve's, you probably want to unblock scripting; right? You want to - you block scripts?

Steve: No, not for security. Keep up your shields for security. I just, you know, I enable it when I need to.

Leo: Yeah. And maybe you can't trust me. And so if you can't trust me, you shouldn't visit my website.

Steve: No, no, it's not about trusting you. It's the by default, because you never know where you're going to go. You click on something, and they're loading scripts from all over the place. It's like, eh, no, thank you.

Leo: Yeah, I know.

Steve: Okay, well, we beat this, we've beaten this to death, my friend.

Leo: It's a great subject, and we've talked about it before. And I know that I annoy everybody with my point of view.

Steve: No, it's important. And I look at the technology and the fact that a meg of download got saved by not going to what Mozilla considers tracking, trackers. So, and I'm happy to see the ads. I do recognize that, I mean, I want these companies to stay in business. It was a blog from Ars Technica, I think, years ago, where the guy said, look, you know, we have a problem because we have a technical audience, exactly as you were saying, Leo. And we need you to see our ads.

Leo: Right.

Steve: Otherwise we can't do this.

Leo: Right. Sometimes I'm tempted just to get out of the tech business, to be honest with you. You guys are hard. You make this stuff hard. Thank you, Steve.

Steve: My pleasure.

Leo: I appreciate it. You can watch this show every Tuesday, 'round about 1:30 p.m. Pacific, 4:30 Eastern - we've got a great schedule on the website, you know, you can see it in your own time zone - 20:30 UTC, live dot - I'm sorry. It's not live.twit.tv. It's TWiT.tv/live. But you can also get on-demand versions at TWiT.tv/sn. We worked hard, by the way, to get those redirects working. TWiT.tv/sn or wherever you get your shows because, you know what, this show's been around so darn long. Episode 512, you know.

Steve: Yup, I know, 2^9.

Leo: 2^9 episodes that pretty much anything that has any claim to offering podcasts

will have it. Or you can always go to Steve's site, GRC.com. He's got 16Kb versions there, the full audio bandwidth. He's got transcripts, a great place to get it, GRC.com. While you're there, pick up SpinRite, the world's best hard drive maintenance and recovery utility. It even works with SSDs.

Steve: Does.

Leo: And check out SQRL, which could be the elimination of all pain.

Steve: We'll see. Yesterday I finished the last major piece of it.

Leo: When are we going to do the SQRL show?

Steve: And they're testing it. And I've got to do a couple more things. Then it's time to go kill bugs and fix things like default pushbuttons and the focus, having it be on the right control. And then we're done.

Leo: Nice.

Steve: So, getting there.

Leo: Very good.

Steve: Close.

Leo: Yay.

Steve: And then back to 6.1.

Leo: If you've got questions, Steve's got answers. Go to GRC.com/feedback, best way, or tweet him, @SGgrc.

Steve: And tips for The New Screen Savers show.

Leo: Oh, yeah, yeah.

Steve: Things that I could talk about that people think would make a nice little one or two-minute segment.

Leo: We need Steve. Steve represents the hardcore arm of the TWiT Army. The hardcore battalion.

Steve: The tech anchor.

Leo: Yeah. Basically you're SEAL Team 6. So we're glad we have you in the TWiT Army. Thanks, Steve. We'll see you all next time on Security Now!.

Steve: Thanks, Leo.