"Topics" Arrives

Description: Today, we have a birthday to celebrate. And then I wound up encountering so many interesting thoughts shared by our terrific listeners that, once I had written everything that I wanted to say regarding the emergence of Google's long-awaited Topics system to replace tracking while still giving advertisers what they need, I'd filled up 18 pages of show notes and run out of space for other news. So next week I'll catch up with everything else that's been happening. But the topic of Topics is, I think, important enough to have most of a podcast for itself.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-935.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-935-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to talk about Google's final, I guess, proposal for advertising without onerous tracking. It's called Topics, and Steve says it's a good thing. We'll also talk about Password rules. Sometimes they can get ridiculous. And then Steve has some very good news for all of us. Stay tuned. You'll be celebrating next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 935, recorded August 15th, 2023: "Topics" Arrives.

It's time for Security Now!, the show where we cover the latest in security breaches, what's going on with MOVEit ransomware with this guy right here, Steve Gibson. Hello, Steve.

Steve Gibson: All that good stuff.

Leo: It's the MOVEit show these days, I swear to god.

Steve: Yeah, for the middle of August.

Leo: Yes.

Steve: And, okay. So there's a lot going on. Today's topic is Topics, which has officially arrived from Google at long last. But we're going to start by celebrating a birthday.

Leo: Oh.

Steve: It's not only 25 years ago that the iMac was created, but it's this podcast's birthday, as well.

Leo: Not quite so old.

Steve: Not quite so old, no.

Leo: Almost. Close; right?

Steve: So, yes. So we're going to talk about that. Then I wound up encountering so many interesting thoughts shared by our terrific listeners that, once I had written everything that I wanted to say regarding the emergence of Google's long-awaited Topics system, which replaces tracking, while still giving advertisers what they need, I had already filled up 18 pages of show notes, and I ran out of space for any other news. So next week I'll catch up with everything else that's been happening. But the topic of Topics is, I think, important enough to have most of a podcast for itself in any event. And so we're going to dig into everything about that. And of course we've got a terrific Picture of the Week. So I think another great podcast for our listeners. And as you can see, another accident here.

Leo: Oh, my goodness. You've got to be careful with those straight-edge razors; you know? You could really get in trouble.

Steve: Yeah, it was so dumb because I had just finished a day of working on SpinRite. And so I showered, and then I shaved. And my mind was a thousand miles away, thinking about some detail of SpinRite, I wasn't even paying attention; you know? And so that's a lesson is like, Gibson, it's only going to, you know, give it three minutes. That's all it'll take.

Leo: I have something, too, after your Picture of the Week.

Steve: Oh, that's right. You said you had a physical...

Leo: I have a little physical thing, a little Picture of the Week. If you can tell me what this is...

Steve: An IoT birdhouse?

Leo: Well, in a manner of speaking. It certainly looks like a birdhouse. There's no room for birds in this thing, though.

Steve: Is there a little red light on the...

Leo: There is a little red light going on and off on there; isn't there. I wonder what that is. Well, we'll show you in just a little bit. I think you'll be very - you'll be, I think, pleased to see what it is.

Steve: Okay.

Leo: All right, Steve Gibson. I am prepared for a Picture of the Week, if you would like.

Steve: So this one is another one that's going to require a bit of description, which I think I can give it. It's a great picture if you just see the show notes. You'll go, ah, that's great, just as you did when you saw them, Leo. So imagine - this is apparently a hotel which has a very vertical profile. And I think on the outward facing edge or side, you know, face of the hotel, is probably a staircase which is enclosed, but at each of the hotel's floors where there's like a plateau between the stairs going up or down to the next level, is an opening. And this opening would be square except that the lower half is cut in a ways. Anyway, the overall effect is that the opening at each floor has exactly the profile of an RJ45 Ethernet jack, you know, Ethernet socket. And somebody looking at the hotel said, you know, that looks like...

Leo: Hey, it looks like my powered switch.

Steve: That's exactly right. That looks familiar; you know? Where have I seen this before? Oh, it's like on every router and every hub and so forth.

Leo: Really do.

Steve: So this person very cleverly positioned themselves in the right place, got a - it looks like maybe a four or five-inch RJ45 Ethernet, you know, wired Ethernet jumper, and is holding it close to the lens with the hotel in the distance, such that the Ethernet connectors are the same size in this perspective as the openings in the side of the hotel, looking exactly like you could plug this cable into the side of the hotel.

Leo: Lovely. Lovely.

Steve: And of course the caption is "We just need a jumper here." Because that's what the...

Leo: That's the short one is a jumper. I love it.

Steve: Yes. A little jumper, a little Ethernet jumper.

Leo: Now...

Steve: Okay.

Leo: I have one for you. If you've listened to this show, and especially if you've listened to the holiday episodes, you know that Steve in his youth created something called the Portable Dog Killer; right? Which was intended really not to kill dogs, but to chase them away by playing high-pitched sounds that only the dog could hear; right?

Steve: Yup.

Leo: Meet the Portable Dog Killer from PetSafe. This is - the reason it looks like a bird feeder or a bird house is, if you've got a neighbor with a dog that bothers you, just as that dog used to bother little Steve, you hang this on the tree right by the fence. When you turn it on, it's got a microphone up here that listens for barking. And when it hears it - and by the way, that's why the red light's going off. I don't know, maybe we're - I don't know if we broadcast wide enough frequencies. People, if you're hearing a very high-pitched tone right now, I apologize. That's the - it's called the PetSafe OBC-1000. See, I think the PDK would have been better. But we have it here because Burke, as you might know, has his little dog Lily, and she barks a lot. So Mikah told him about this and said, if you buy this and put it in the hall...

Steve: We will all thank you.

Leo: No, I love it. Lily can bark all she wants. Doesn't bother me. But what I thought was interesting is you thought of this years ago, never commercialized it. But here it is as a commercial product.

Steve: Very cool.

Leo: The PetSafe Portable Dog Killer in the form of a birdhouse. Isn't that amazing? See, you were ahead of your time.

Steve: Again, I was, what, 15 or 16. So to me, calling it the Portable Dog Killer, I mean, and it looked like a laser gun. So, you know, and we were all - we were raised on "Lost in Space" and "Star Trek." And so, you know, that's what I was going to call it.

Leo: I think this is commercially probably a little more acceptable.

Steve: I think so, too.

Leo: Just, yeah, keep it arm's length away from human ears to avoid hearing damage. If a dog begins barking while you're setting it up, mounting or hanging the - it's called the Outdoor Bark Control, OBC-1000.

Steve: Ah, there we go.

Leo: Yes. But Burke came running in and said, you should turn that off if you're going to put it on the air. It might deafen our listeners.

Steve: Okay. So Leo, you and I recorded Episode 1 of Security Now!...

Leo: Oh, my.

Steve: ...on August 19th of 2005.

Leo: Oh my, Wow.

Steve: Today is August 15th of 2023.

Leo: Holy cow.

Steve: Uh-huh. Which means that with today's podcast number 935, we will have finished our 18th year. And next week's podcast will be the start of our 19th year. And when I went back to check the date, you've got it on the screen, of that first podcast - which was, by the way, all of 18 minutes long - I think probably I just...

Leo: You want to hear a little bit of it?

Steve: I just took a little stop between...

Leo: I think it sounds exactly the same. I'll be honest. Oh, no, we didn't have our music yet. This was some other music I was using.

Steve: Oh, that's very laid back, yeah.

Leo: "This is Leo Laporte, and I'd like to introduce a brand new podcast to the TWiT lineup..."

Steve: Aww.

Leo: "...Security Now! with Steve Gibson. This is Episode 1 for August 18th, 2005." I like the music. "You all know Steve Gibson. Of course he's on TWiT regularly, This Week in Tech. We've known him for a long time. He's been a regular on The Screen Savers and Call for Help. And, you know, he's well known to computer users everywhere for his products. He's very well known to consumers for" - I used to give you a better intro - "SpinRite, which was the inspiration for Norton Disk Doctor and still runs rings around it, is the ultimate hard drive diagnostic, recovery, and file-saving tool." It still is. I say that still every week, except now I say "mass storage."

"He's also been a very active consumer advocate, working really hard to help folks..."

I think my voice sounds higher, but I don't think it actually is. I think it's a little bit of a bit slip in this thing. Let me play a little bit...

Steve: It's funny because I do, hearing your voice, I remember that Leo.

Leo: Yeah, the young Leo. And you're on Skype, so the quality of your audio is terrible, even on Screen Savers.

[Episode 1 continues playing in the background]

You don't sound different, but it sounds phased a little bit. We got you a PR 40, more bandwidth, started using Zoom.

Steve: Yeah. And I was on the other side of that T-1 line, so I didn't have...

Leo: Oh, you had 1.44Mb.

Steve: I didn't have today's bandwidth.

Leo: There you go. That makes a big difference. If people want to hear this, it's still on the website, and the shortcut is TWiT.tv/sn1. Sn1, the first Security Now!.

Steve: So that episode was titled "As the Worm Turns."

Leo: Yeah.

Steve: The first Internet worms of 2005.

Leo: Holy cow.

Steve: I know. And its description made me shake my head because it reads: "How a never-disclosed Windows vulnerability was quickly reverse engineered from the patches to fix it and turned into more than 12 potent and damaging Internet worms in three days."

Leo: Wow.

Steve: What does this mean for the future of Internet security? And here we are, having just, you know, celebrating the 18th birthday of this podcast, and so much has changed, and so much has not. So anyway, thanks to feedback from our amazing listeners, one of the things that's been driven home for me during the past 12 years is how much this

podcast means to our listeners and, I suppose, how much it would be missed, at least for a while, if it were to ever end. Now, obviously, it's going to end sometime. Unfortunately, Leo...

Leo: Unfortunately, in 65 episodes.

Steve: We're not both going to live forever. When William Shatner, who is currently 92...

Leo: Who is going to live forever, I believe, yes.

Steve: I think apparently. But, you know, he's in remarkable physical and mental health, and of course he recently took that quite emotional for him ride into orbit and back. He was asked about his secret to long life. And he replied simply, "Don't die!" Right? That's it. Don't die.

Leo: Yeah.

Steve: So I'll confess that while my middle name is not Tiberius, I'm going to do everything to follow the Shatner plan. And thanks to our listeners, it occurs to me that perhaps this podcast should follow the Shatner plan, too. As all of our listeners know, I've been talking about ending my involvement with Episode 999. I'm here after 18 years, Leo, due to our gentleman's agreement to do a podcast together. And I think that I should remain here as long as that's what everyone wants.

Leo: What? Oh, my god. You just made, not only me and everybody in this building, but about 100,000 listeners extremely happy. You're saying you're willing to go beyond 999?

Steve: So, yes.

Leo: How are we going to get the four digits? Can we start over at zero?

Steve: I've written some code in my life. I can figure out how to change three digits into four.

Leo: Steve. Be still my heart. Oh, my gosh. Thank you, on behalf of everybody listening.

Steve: So it seems to be the case that even after 18 years, everybody still wants this.

Leo: Yes. Yes.

Steve: And, you know, and I feel as though we have a lot of leverage, by which I mean that this podcast appears to matter to a lot of people.

Leo: Yeah, that's true.

Steve: And that's enough for me.

Leo: Lisa just came running in. What? Let me see if I can find a shot with her in - well, you'll have to come around here and do it. She's very happy.

LISA: Yay.

Leo: Yay.

LISA: And we still need to come down and harass you. Now I'll stop because everyone will get mad. But I was like...

Leo: There's a lot of excitement, and I can hear the cheers in the other room.

LISA: Oh, everyone was freaking out.

Leo: Patrick Delahanty's jumping up and down. Yeah, that's really great. So by the way, Patrick just told me, he's updated TWiT's code, in fact he did it a couple years ago, to support more than 999 podcasts. So Steve, if you're willing to go 999, I'll do it, too.

Steve: Well, you know, Leo, these first 935, they just flew by. You know? It's like...

Leo: I was really feeling like you were kind of getting tired of doing it, to be honest with you. But you're not.

Steve: Well, certainly as we head into year 19, I think we've at least established that we're not about to run out of material.

Leo: Yeah. By the way, in the description of Security Now! 1, it says "this short podcast." It literally, I thought at the time that it would be a short podcast. But bad guys just don't rest. Hey, thank you, Steve, and congratulations. That's really, really good news.

Steve: Well, it feels - I've known for some time, and I thought, this is the occasion where, on our 18th birthday, that I should just say, you know, why stop doing a good thing?

Leo: Oh, thank you, Steve. On behalf of the entire Internet community, thank you.

Steve: So a bunch of Closing the Loop feedback from our listeners. Some guy whose handle is gimix, G-I-M-I-X, and then cubed, Gimix^3, his actual name appears to be Jordi, he said: "Hey Steve, today" - and this was on Friday - "Chrome greeted me with this dialog." And I have a picture of it in the show notes. And it says: "Turn on an ad privacy feature." And it goes into some details.

And so he writes: "Google adding a privacy feature? And asking me PERMISSION to turn it on? Aha. Extremely suspicious. I guess that's FLoC, which was implemented anyway, even if the community pushed back. Any advice on whether we should turn it on or not? Maybe material for the next episode! Greetings from Barcelona, longtime listener, Jordi."

I received Jordi's note, as I said, on Friday, and I knew what it was. This was not FLoC, this was the long-promised rollout of Topics, which is Google's replacement for hidden tracking-based profiling in favor of transparent site-based interest profiling, which is performed entirely client-side, in TNO style, and is entirely under the control of each individual Chrome user. I mean, it's a breakthrough. But I was greeted by it when I opened Chrome yesterday. Since I'm no longer a Chrome frequent flier, I hadn't seen the dialog when Jordi had tweeted the picture of it to me.

Since the Chromium browser's Topics API is finally making its appearance, and since there's reason to believe that this will be the system that changes the world, it's time to revisit what it is and how it works, which we're going to do shortly. So anyway, Jordi, thank you for being the first to point it out. I got the dialog of my own, and everybody else who's using Chrome will have been receiving this. So we're going to talk about exactly what this is and how it works.

Leo: And why you should use Firefox.

Steve: And why you should use Firefox. Well, actually, I'm hoping Firefox will adopt this. I sincerely am.

Leo: Oh, okay. That's a twist. Okay.

Steve: Oh, yeah, yeah. I think - and you and I are going to have some fun with this, Leo, because there are, you know, multiple sides to this.

Okay, now, someone who chose the unfortunate handle "burnedeye," I don't know why, sent something very interesting. So burnedeye tweeted: "Hi, Steve. You may find this interesting. I know you're a Firefox user. So am I. Recently I have found out about this fantastic extension called 'Firefox Multi-Account Containers.' It solves an issue of being automatically logged into, for example, all Google services, when you only want one, for example, YouTube.

"You can isolate YouTube.com into its own container, where you can be signed in while all the other Google sites such as Google.com, Gmail, et cetera, can stay in a default container where you're not signed in, and those sites won't be affected by the fact that you're signed into YouTube." He says: "It's a browsing tab virtualization, in a way." He said: "Love it, just as I love the Security Now! podcast. Longtime fan, Tim."

Okay. So I was not aware of this slick browser extension, and it does solve a problem I sometimes have. Its author explains thus. He says: "The Firefox Multi-Account Containers extension lets you carve out a separate box for each of your online lives. No more opening a different browser just to check your work email. Under the hood, it separates website storage into tab-specific Containers. Cookies downloaded by one Container are not available to other Containers. You can even integrate individual Containers with Mozilla's VPN to protect your browsing end location.

"With the Firefox Multi-Account Containers extension, you can sign into two different accounts on the same site, for example, you could sign into your work email and home email in two different Container tabs, even if they use the same server; keep different kinds of browsing far away from each other, for example, you might use one Container tab for managing your Checking Account and a different Container tab for searching for new songs by your favorite band; avoid leaving social network footprints all over the web," he says, "for example, you could use a Container tab for signing into a social network, and use a different tab for visiting online news sites, keeping your social identity separate from tracking scripts on news sites; and protect your browsing activity in individual Containers using Mozilla VPN, so you can shop while traveling abroad, but check your bank account from a server in your home country."

And finally: "After installing the Firefox Multi-Account Containers extension, click the Containers icon to edit your Containers. Change their colors, names, icons. Long-click the new tab button to open a new Container tab."

Anyway, okay. Though I haven't tried it, this seems very cool. The one thing I would caution is that he writes "keeping your social identity separate from tracking scripts on news sites." But as we know, and as we saw quite vividly last week with the lengths Yahoo!, in that example, goes to in order to track people, containerizing explicit authentication cookies is not the same as fully anonymizing one's appearance on the web. I'm certain that someone examining the queries being admitted by, or JavaScript running in, adjacent containers would be able to detect that they're running side-by-side in the same browser. Nothing is simpler than noticing that the queries are all emerging onto the Internet from the same private IP address.

You know, he addresses this simple IP-based tracking with the feature of an automatic tie-in with Mozilla's VPN service. But this is not to say that the idea of being able to be logged into the same service under multiple identities with a single browser is not extremely useful. I often, you know, wish I had a different account at the same place. And I do tend to use a different browser if I need to. I remember when I was messing around with SQRL a lot, that was something I was having to do. So anyway, I wanted to share Tim's tweet so that we all know about it. And those of us who Firefox users are able to take advantage of it.

Matthew Dudek, he said: "Hi again. Have a question about Full Disk Encryption on SSDs. If an SSD is already in use in an unencrypted state, is it then impossible to fully encrypt it with BitLocker or VeraCrypt due to data stored in inaccessible blocks because of over-provisioning and wear-leveling? How can one encrypt an in-use SSD and guarantee all data is sufficiently scrambled? Thanks again!"

Okay. So since I plan to fork the early work on SpinRite 7 into a separate product named Beyond Recall, I will eventually acquire a great deal of firsthand experience solving these problems. Matthew is correct in assuming that wear leveling, defective region sparing, and over-provisioning inherently take data-containing mass storage regions out of service, rendering them inaccessible through the mass storage device's normal data API. There are two different means for fully erasing all traces of data, even data that's inaccessible. The problem is that they erase all traces of data.

What Matthew wants to do is to take an SSD that's already seen some use and add external full disk encryption to that existing device while not leaving the inaccessible regions, which might still contain some previously in-use unencrypted data, unencrypted. And as far as I know, that's not possible. He'd like there to be some command to destructively and permanently erase only all of the inaccessible regions. But I'm not aware that any such command exists. This means that he would need to copy the drive's contents to another device, then arrange to securely erase the entire drive, which would and does include all user-inaccessible areas, then implement BitLocker or VeraCrypt or whatever encryption on that drive, then restore the drive's original contents.

So unfortunately, no way I know of, of just doing it in place. That inaccessible previously unencrypted data will stay there until you perform a secure erase on the drive, which is the only thing that will get rid of it because you can't get to it otherwise. Most of these storage controllers have manufacturer proprietary, undocumented, typically unknown, you know, backdoors that allow you to manipulate stuff. That's something I may end up getting into, depending upon how things go with Beyond Recall. But, you know, nothing known publicly.

Trevor Welch said: "Hi, Steve. I have a SpinRite question. I have a mixed array of drives, hard drives, SATA SSDs, M.2 SSDs, and probably even some weird proprietary external hard drives. I'm finally getting my act together and going to be backing them up to a NAS and backing that up to the cloud. Is there any advantage to running SpinRite on any of these drives before I do this? Or do I run the risk of maybe pushing one of the drives into its demise while running SpinRite, and being unable to then copy all the information off? Some of these disks are very old, 10-15 years maybe, and are in unknown condition as of right now. So I just want to make sure I give myself the best chance of being able to get as much data as I can off of them." And he says: "Love the show and excited for SpinRite 6.1 to come out. Maybe an announcement on Tuesday?"

Well, okay. No announcement today. As always with a project of this size, with as many moving pieces as this has, working to get to the "there's nothing left to be done" state reveals additional things that need to be done. And I'm reminded of that old thought puzzle which suggests that it's impossible to actually get to your destination because before you can, you first need to get halfway there. Then halfway there again, and halfway again, and halfway again, and so on indefinitely, thus theoretically unable to ever reach the goal. But, you know, we are down, the good news is, to very few remaining known things that need addressing. So yeah, it's looking like one of these weeks very soon I will have an announcement for the listeners of this podcast.

As for Trevor's question, I don't see any reason to run SpinRite on any of those drives ahead of time. Just try copying all of their data off at the file level, you know, file by file. Copying programs are notoriously finicky about hitting errors during their work, but there are copying utilities that will retry for a while, then skip over any trouble they encounter. Even the old XCOPY program from back in the MS-DOS day, which is still present, it has an "ignore errors during copying" option. And I'm a fan of Robocopy on Windows. It's got all kinds of extra feature switches that allow this to happen.

The one thing those tools won't and cannot do is deal with problems in the file system's metadata, which SpinRite can do, since they depend upon that metadata to find the file names and the file content locations in order to copy it. But otherwise, it's usually possible to get most of the data from a drive. But if any trouble is encountered along the way, I'd say by all means let SpinRite have the drive to see whether it's able to fix those areas that may have been troublesome.

Matt G, tweeting from @mpgagne, he wrote: "This might be the most annoying password rule list I have ever seen."

Leo: Oh, this is a game, I bet. Is this the game?

Steve: Well, this is really - no. This is really interesting because it's got some hidden gotchas. He said: "You can't use a password generator because any repeat of a single character makes the password invalid." He said: "Thought you would enjoy."

Leo: Oh, please. That's terrible.

Steve: Okay, now, actually this is - there are so many good things here. So, okay. Here's what the rule list is. And by the way, as we'll see from the bottom rule, this is JPMorganChase's password creation guidance. So first, must be 8-32 characters long. Must include at least one uppercase, one lowercase, and one number. Okay. Those are kind of standard. Must not have special characters or punctuation. Okay.

Leo: I don't like that.

Steve: Get this. Must be different than your previous 24 passwords.

Leo: What? Where did that number come from? What?

Steve: I know. Oh, it gets better, Leo. Must not include your email ID partly or fully.

Leo: Okay.

Steve: Must not include your first name or last name.

Leo: That's fair. All right.

Steve: Okay. What if your last name or first name has, well, okay.

Leo: Is monkey? What are you going to do then? Huh, huh, huh?

Steve: Here's a problem. Must not include more than 2 identical characters.

Leo: That makes no sense.

Steve: Okay. Get this. Must not include more than 2 consecutive characters.

Leo: Wait, what?

Steve: I know. I know. And you can't have more than 2 consecutive characters. Okay. Must not use the name of the financial institution - JPM, MORGAN, JPMORGAN, CHASE, JPMORGANCHASE, or JPMC.

Leo: I wonder how much those last couple, not the last last one, but the consecutive characters and identical characters, reduces the password space? That must cut it way down.

Steve: Oh, Leo, you have been paying attention to this podcast. I credit you that, definitely. Matt's tweet included a screenshot of those requirements, which we just shared.

Now, there are three problems that come to mind. First, one of the rules reads "Must not include more than 2 consecutive characters." Yet we know from the first rule that the minimum password length is 8 characters. So it's unclear how you create any password longer than two characters if you must not include more than two consecutive characters.

Leo: Well, you have to put a number in between them; right? They mean alphabetic characters?

Steve: It's not clear. They just said "characters."

Leo: Characters.

Steve: Yeah. This guy was a character, whoever wrote this.

Leo: Yeah.

Steve: Now, so it must be that the author of this rule meant to say "2 consecutive identical characters."

Leo: Oh.

Steve: Right?

Leo: Maybe.

Steve: Except that the preceding rule is...

Leo: You can't have any identical characters at all.

Steve: But no, it says: "Must not include more than 2 identical characters." Now, again, here the digit "2" is not necessary because you could just say "identical characters." Right? Because it's redundant. But okay, fine. So that means that the misworded following rule, the "must not include more than 2 identical characters" is redundant because it's fully covered by the one that precedes it.

Okay. But aside from that grammatical nitpicking, there are two bigger problems. The very first rule states that the password must be 8-32 characters long. Okay. So a minimum of eight characters? Really? An eight-character password is sufficient, given all of the rest of the rigmarole that JPMorgan Chase customers are being put through?

Leo: Nonsense. Just nonsense.

Steve: And when you think about it, Leo, given how difficult they've made it to create any password...

Leo: Yeah, because your password manager doesn't know these rules.

Steve: No, no. So, no, it's got to be done by hand. But given how difficult they've made it to create any password that somehow manages to get through the gauntlet of those rules, users would be hugely incentivized to quit after somehow working out an eight-character string that qualifies.

Leo: Yeah, you're right. You're guaranteeing an eight-character password.

Steve: That is exactly right. And that brings us to the third and worst problem of all. These ridiculously - exactly as you immediately saw - these ridiculously onerous rules are going to drive users to create the shortest possible passwords while at the same time making brute force guessing vastly easier and more practical. Any intelligent brute forcer will be informed by the same limiting rules as the password creators. So this dramatically and incredibly reduces the possible brute force search space. You know, "No special characters or punctuation?" Wow.

Leo: I always hate it when I see that one. I hate that.

Steve: Yes, that's so stupid. Talk about dramatically reducing the alphabet size and thus the search space. Those same rules which make it difficult for a customer to create a qualifying password, automatically discards a vast universe of passwords that would have been possible, and that an attacker would have needed to try. But now the attacker already knows that those would not have qualified.

Leo: Oh, don't worry about those. Yeah, that's great.

Steve: Imagine all of the guesses where more than two characters are the same. None of those ever need to be tried.

Leo: Unbelievable.

Steve: What a perfect real-world example of someone thinking that they're being quite clever by forcing their customers into compliance, when they're inconveniencing those customers, while at the same time making things far easier for the attackers. Wow.

Leo: Unbelievable. And we see this all the time.

Steve: At this point, yeah, it's just, you know, trying to make it better and making it worse.

Leo: Go ahead. You want to take a break?

Steve: I do, yeah.

Leo: I just want to know if you've seen the Password Game. Because this is so fun, from Neal.fun. Have you ever heard of this? Okay.

Steve: I haven't, no.

Leo: Okay. I'm not going to show you this if you have anything to do tonight.

Steve: Uh-huh. Whoa.

Leo: Okay. So it's at Neal.fun.

Steve: Maybe don't show it to me until after SpinRite 6.1 ships.

Leo: It's to make fun of all of these. Okay. So I'm going to do monkey. It has to include a number, 123. Okay. Uppercase letter, okay, I'll make the "m" an uppercase.

Steve: Oh, yeah. That's good.

Leo: The digits in your passwords must add up to 25. So 3 plus 8 is 11, plus 9 is 20, and then a 5. Okay. Password must include a month of the year. June. Okay. One of our sponsors, Shell. Roman numerals should multiple to 35. Uh-oh. So this has to be maybe - so that's seven times five. So maybe this will be VII, and then this will be V. Okay. Oh. CAPTCHA. D22BDD22EZ.

Steve: Now, does it always give you the same requirements? Or do those...

Leo: Uh-oh. Yes. So I just ruined it by doing the CAPTCHA because there was a 22 in there. So now I can eliminate everything but a 3. Okay. It's 22. So it's 2 plus 2 is 4. Okay. I get it. 13, 20, 5. Okay. Here's where I stopped. Your password must include today's Wordle answer. By the way, it is today's Wordle answer. I've done this before. John and I have both done this. So you now go off to solve Wordle.

Steve: It checks?

Leo: It checks. So you now go off to solve Wordle. It will give you a chess problem in just a bit. At one point your password starts to catch on fire, and you have to [indiscernible].

Steve: Anyway, for our listeners, it's Neal, N-E-A-L, dot F-U-N.

Leo: Yes, Password Game. He does a lot of fun games. I guarantee you, because you're smart, you will dig this. I have never gotten past the fire. John says he's gotten past the fire. What's the last step you got to, John? How far did you get? I think he said he got 25 rules, something like that. It goes on. The guy who created it, Neal, said he's never finished it. You will enjoy it, Steve.

Steve: Wow.

Leo: All right. Now we're going to take a break. Please, folks, do not go off and do that. Stay tuned. You can do it tonight. Enjoy the rest of Security Now! before you engage in this. It's like the paper clip game. Once you start, you know, you kiss the rest of the day goodbye.

Steve: Anybody who's ever looked at underpasses being built in California has noticed that they are no longer brittle. They are resilient.

Leo: That's right.

Steve: They are something sitting in a pocket so that if an earthquake shakes it, it doesn't crumble, it just rocks around a little bit and hopefully returns close to where it started.

Leo: What's the worst kind of building to be in in an earthquake? It's not a wood frame building. It's a brick building. Those are deadly, yeah.

Steve: No give.

Leo: No give.

Steve: So Josh Randall, he said: "In SN-934, you mention that many sites now require you to create an account with an email so that they can spam you later. That's precisely why I use DuckDuckGo's email alias - that's duckduckgo.com/email - which I recall you were rather negative about a few episodes back. When I create an account with any new service or site now, I let DDG create a new, random email address for me that is an alias to my main DDG account, which is in turn an alias to my actual email. I can receive messages at my actual email through any of these aliased DDG emails, and reply to those messages without ever revealing my actual email. And if any site or service ends up spamming me through one of my random aliases, I can simply deactivate it and, poof, no more spam."

Okay. So let me just be clear about my negativity. I wasn't aimed at DuckDuckGo specifically. I'm 100% behind the use of email aliases. I think, as the kids would say, Leo, they rock, and I use them myself all the time. I have hundreds of email aliases. But the thing that's different is that the aliases I use are created by my own email server at GRC.com, which I'm not going to decide to suddenly terminate. Or if I were to, it would be my decision, under my control. Given the importance of email as our account recovery and proof of identity, I would be nervous to be asking any third-party provider for such a service if I used it in such a widespread fashion.

Unlike a use-it-once credit card number, an email address needs to be inherently static and persistent. If DuckDuckGo were ever to decide to terminate that service, it would create a significant inconvenience for its users who are using it in this fashion, who would need to manually change every one of their registered email addresses everywhere they had ever used them. Again, I love the idea, and I get it that not everyone is able to run their own email server, especially since consumer ISPs actively block their customers from running local email transports. But anyway, so love the idea. And I guess I wish there were a way to do it that seemed really safe.

Leo: Yeah. I mean, I never even thought about that. I have my own, you know, at Fastmail I have my own domains, so I can have an infinite number of @domains. So I just do that. And those are all unique. But maybe a little bit more guessable than, you know, xyz@zmzdd.

Steve: Well, and as you know, because you've been following me for a while, my email address changes annually.

Leo: Yeah, you're smart. That's such a good idea, yeah. Because you don't want anybody to reach you.

Steve: Right. Or, you know, they can on Twitter or by talking to Sue or Greg.

Leo: If you know the algorithm, if you are a friend of Steve's you will know the algorithm, and then you can reach him.

Steve: David Halliday, he said: "In SN-933 you highlight that Russians are now prohibited from contributing to open source. However, any GPL product makes it very clear that the user MUST contribute to the project for the license to be valid. Thus the new Russian Astra Linux-based OS cannot be legally possible."

Leo: Oh.

Steve: Okay, now, as I understand it, the GPL requires that any improvements which are made, thanks to having had access to the source code, must be returned to the project. And so, yes, I think that David's point is correct, not that Russia will be particularly concerned about violating the licenses of the West. Russia has essentially stolen Linux from the Linux project; but with the rising political tensions, who's surprised by that?

Ooh, some math. Thomas Apalenek, hope I didn't mangle your last name too bad, A-P-A-L-E-N-E-K. He said: "Re: Satellite Crowding. Steve. I really enjoyed the episodes and satellite hacking and related satellite info. However, the follow-up discussions on swarms of satellites needs to be put in perspective." For which we thank Thomas. "A 3D graphic of white dots showing all the satellites and debris orbiting the earth does look a little bit frightening. However, if the satellite size were displayed at the correct scale relative to the earth on the graph, you wouldn't actually see any satellites at all, except for possibly the International Space Station. Starlink's 40,000-plus satellites, in particular, sounds like a lot. But if you imagine 40,000 cars equally spaced over the entire surface of the Earth, it doesn't seem nearly so bad.

"The effective surface area of a 300km orbital shell is about 560 million square kilometers. An 1,100km orbital shell is about 700 million square kilometers. The Starlink orbital plan includes three shells ranging from 340km to about 1,100km. The debris issue is a concern, and launch agencies do need to make sure it doesn't get out of control. For now, all of the items in orbit are, on average, thousands of miles away from each other. The skies won't be darkening anytime soon." And again, Thomas, thank you for that very valuable perspective. That's very useful.

Leo: That's a good, excellent point. Space is big.

Steve: Space is big.

Leo: Really big.

Steve: Yes. Yes. Which is why, you know, things are zooming around and mostly not hitting each other.

Leo: Right.

Steve: When they do of course it's pretty spectacular. But anyway. Brian Norwood said: "Catching up on this week's episode, and this is what came to mind on the Voyager 2 segment with the shout." He said: "Next year we'll find out we let a Borg equivalent know exactly where we are." Remember that NASA was able to regain control.

Leo: [Crosstalk].

Steve: Well, he said, remember, NASA put out a high-power pulse in order to get Voyager 2 to reorient itself. Now, I think I've noted before as a sci-fi enthusiast looking at all the trouble we have right here on Mother Earth, where we're all basically variations

on humanity, that as I grow older, and hopefully somewhat wiser, I am coming to appreciate where I once used to bemoan how difficult it appears to be to travel between the stars, there's that famous and pithy observation that "Good fences make good neighbors." So these days I'm much more satisfied with reading fantastical stories and books than actually having any first contact. You know, let's hope they're not coming here to take our water because that would not be good.

Brian Gluck said: "Can you please share the name of the session manager that you use in Firefox," which I mentioned last week. He said: "I had one that I loved that saved all of the open windows and tabs, but it stopped working a number of years ago. I think it was called Session Restore, but I honestly don't recall. I would love to know which one you're using. I miss this functionality and don't like to download an add-on into Firefox without a recommendation from someone I trust."

Okay. It's called "Tab Session Manager" as three words, and I just noticed that it's also available for Chrome and Edge. And it really is a nice piece of work. I've been using it for many years, so I'm happy to vouch for its value and stability. It's free and user supported. And having just written all that, I just sent its author $25 dollars through PayPal since I would like to keep it around. It is, you know, it's open source, it's on GitHub, the guy's maintaining it, and Firefox does sometimes make a radical change that requires rewrites of their add-ons, which is exactly how Brian's earlier thing, that Session Restore, you know, some guy created it and wandered away. Then Firefox changed and broke the extension. So the fact that this thing is being maintained is a good thing.

And finally, Rusty tweeted: "@SGgrc How many people have pointed out that the 2 degrees off that Voyager was pointing at 12.3 billion miles was a bit under 430 million miles, or about five times the average distance between the Earth and the Sun." He says: "It's fantastic that they were both able to hear and send."

And so, Rusty, you take the prize as the only listener whose tweet I've seen took the time to do the math and give us a calculated answer. My own intuition suggested that this had to be the case, but it's nice to have some numbers to go along with it. So about five times the distance between the Earth and Sun is how far off angle, away from us, that beam was passing when Voyager was trying to talk to us. We weren't hearing anything that it was sending any longer.

Okay. And Leo, let's do our last break, just so that we can do Topics in one whole piece.

Leo: I like it. By the way, I'm making some good progress on our password.

Steve: Oh, my goodness.

Leo: I did have to identify this country. And it took me a little while, you know. But the name of the hotel was the giveaway. I looked that up on the Internet. So now I just have to get a leap year in here without changing the addition. Maybe I'd better do the ad instead.

Steve: What is that black blob?

Leo: Oh, that's the current phase of the Moon as an emoji. Yeah, it's a waning crescent right now.

Steve: Oh.

Leo: So it says insert the...

Steve: I did get some crap from kids in school over the "waning Gibbon."

Leo: The waning Gibbon moon, yes.

Steve: That's right.

Leo: Yes, this is the waning Gibbon. Anyway, I'm only, you know, a fraction of the way along. But I've got a great password. I must say.

Steve: Okay, now...

Leo: Does Chase let you use emojis in your password?

Steve: I want to say, though, you missed our first episode where we talked about Topics. And I need to have your attention for this one because it's important for you to get this. I'm just saying.

Leo: Okay. I'm not moving. I had to get up and run up and down the hall and celebrate, and I did miss that first image. But I won't miss this next one. I'm paying attention now. All right, Steve. I am paying attention. I am not leaving this seat. I am listening to every word. Go right ahead.

Steve: I think you're going to find this useful, as our listeners will. So, okay. The right answer doesn't always present itself the first time. When we encountered "FLoC," which should stand for "Failed Learning of Cohorts" instead of "Federated Learning of Cohorts," you explained, Leo, that in referring to it as FLoC, Google was clearly struggling to keep with a bird theme. It turns out they were struggling too hard and the FLoC flew the coop.

Leo: Yes.

Steve: But as I said, the right answer doesn't always present itself the first time. Sometimes it's necessary to experiment and iterate. The reason SpinRite took three years is that there were a lot of avenues I went down that I ended up backing out of. It's like, okay, well, that approach didn't work to be a universal solution, so let's try this. So the design of Topics, which is Google's final solution for the replacement of profiling by tracking, shows that Google has learned a great deal from their previous attempts. And we not only have a system that's ready for prime time, it's beginning to roll out now. It's in Chrome today.

Now, we initially covered the Topics API about a year and a half ago, immediately following Google's first announcement of this proposed new system. This occurred, Leo,

during one of your rare absences from the podcast. Security Now!'s February 1st, 2022 podcast was titled "The Topics API."

Leo: Oh, I know why. I was in Portugal at the time and having a grand time. So there.

Steve: Okay. Well, we're glad for you.

Leo: Yes.

Steve: Jason held down the fort. And I have always regretted that you missed that discussion because having you understand this is critical because, you know, you appear, not surprisingly, on many podcasts here.

Leo: Yes, I do. I'm listening.

Steve: So fortunately, we have another shot at this. And this time it matters even more because, unlike FLoC, Topics is probably going to fundamentally change the way the Internet works.

Leo: Hmm.

Steve: Now, as we noted last week, DNT stood for Do Not Track, whereas GPC (Global Privacy Control) is an explicit request for privacy enforcement. In other words, they're not the same thing. GPC is not about tracking, even though much of the tech press refers to it as an anti-tracking measure. Similarly, Google's Topics solution is not some new means for tracking users on the Internet, even though virtually all of the tech press is calling it that. I think the problem is that since tracking is all we've ever known, and change is difficult, everything is assumed to be some form of tracking.

But as everyone is going to understand by the time we're finished here today, Google's Topics system is explicitly and almost painfully a non-tracking solution. If I may be permitted to use the term, it is truly a privacy-forward system for allowing websites to learn a little something about the topics which may interest their visitors. Period. Full stop. Again, "Topics is a privacy-forward system for allowing websites to learn a little something about the topics which may interest their visitors."

It is a means for allowing Google, you know, the Internet's massive advertising behemoth, to continue to deliver user-relevant advertising without tracking. Topics does not utilize any sort of tracking. None. It's basically Google seeing the writing on the wall that tracking may not be permitted indefinitely, but they'd like to continue to exist indefinitely. They know that it might eventually be outlawed, at least in some jurisdictions. So they want to have some sort of user-profiling replacement ready if that happens. And in fact they've already announced that Chrome will begin deprecating its support for tracking via cookies starting next year, in 2024.

Now, I used the term "user-profiling" just now, deliberately, because while that's what Topics is, it is an entirely different and vastly weaker form of profiling from what we're accustomed to. With profiling via tracking, many hidden entities on the Internet know

everywhere we go. A visit to a site for erectile dysfunction gets logged into many hidden databases over which we have no control, and that information is then made available for sale.

Leo: I didn't do it. Someone else was using my computer.

Steve: That's right. Somebody borrowed your computer. Yeah. You left it unattended at Starbucks.

Leo: Yeah, that's it. It was Lisa. Yeah, that's it, yeah.

Steve: So it's hardly surprising that's not anything that anyone wants to have everywhere you go, everything you do, being secretly logged into hidden databases over which we have no control. Under the Topics system, and assuming an absence of tracking which is a separate issue we'll get to in a minute, only your local browser sees that visit. And that visit is not recorded. It only knows that this is a site having the topic of "Health and Wellness," so that "Health and Wellness" topic gets added to the topics you have shown an interest in for the week. That's it.

Thanks to the creation of this new technology that we're going to go into detail, Topics provides a means by which a user's web browser may learn by inference about its users' current interests, by virtue of where they take their browser on the Internet. And the browser is then able to judiciously make a few of those interests known to websites and advertisers who ask.

Okay, now, before we get into the technology, I want to make it clear that, as usual, we're here to talk about technology, and that I'm deliberately agnostic on matters relating to the non-technical questions of whether or not any profiling is a good thing, whether or not it is driven by tracking. Is any form of profiling okay? I get it. That question is controversial. There are those who feel very strongly that all use of the Internet should be anonymized to every degree possible. Like anybody at the EFF, apparently. They feel that it is their right to minimize the value they present to a website - and its supporting advertisers - by remaining as anonymous and as unknown as possible. They feel that every visit to any website should be siloed, with no other site or third-party content provider having any information about them that they don't supply to the site being visited.

I get that. And, you know, Leo, you and I were just discussing here last week whether the value obtained by user profiling is really worth what advertisers believe it's worth. Now, that's not something I can speak to. Google and advertisers and sites like Yahoo! clearly believe it is. Or perhaps they just want it because that information is available, and they'd like to know who they're spending their advertising money on. Perhaps that feedback allows them to better tune the content that their sites offer. You know, okay, whatever.

What I do know, though, is that user profiling via tracking represents the height of privacy intrusion. As far as I know, an immutable record of every website I have ever visited is squirreled away in multiple massive hidden and inaccessible-to-me profiling databases. And I have zero control over that. That's the world we're in today. But if Topics succeeds, and Google would appear to be in the position to singlehandedly deliver its success, it is a far less intrusive profiling technology. And in addition to being a much weaker information gatherer, Google has chosen to provide its users complete control

over the Topics their browser presents to the world, including turning it off altogether for full anonymity. I'll explain that further in a minute.

So if only on that basis, Topics at least represents a huge step in the right direction. Yes, by default some interest profiling remains. But the means of obtaining those significantly weakened profiles is no longer tracking. And users have complete visibility into their online profile and are able to curate, edit, and even delete any of it or all of it as they choose. So it's a compromise. But there are many websites begging for our support. My feeling is, if voluntarily letting them know something about who we are allows them to generate, as they claim, significantly more revenue from our visit, is that too high a price to pay? Again, it's an individual decision. But now, in a world with Topics, at least, it's one we're able to make.

Another of the arguments presented by the naysayers is that if Topics is embraced we have no guarantee that tracking will end, and that we won't merely be adding another powerful profiling technology on top of the existing mix. And those people are right. But as I noted above, Google is planning to begin the deprecation of explicit tracking support via cookies once Topics has come online. And if Google truly wants to prevent tracking once they no longer need it, neither it nor its bad press, they're in the perfect position to do so. Ultimately, though, as with GPC, it may come down to legislation. And that's fine, if legislation is what's required.

Do Not Track might yet be reborn, this time enforced by local or national legislation making it illegal to track someone who has indicated that they don't wish to be tracked. And the European Union is likely to outlaw it altogether. If a viable alternative to profiling via tracking is present, that is, an alternative to profiling via tracking is present, right, as it will then would be thanks to Topics, it will be difficult for the trackers to make the case for why users need to be tracked to support the Internet's advertisers and content creators.

Okay. So three and a half weeks ago, on July 20th, TechCrunch noted that Topics was finally arriving. They wrote: "Google continues the rollout of its Privacy Sandbox APIs, its replacement for tracking cookies for the online advertising industry. Today, right on schedule and in time for the launch of Chrome 115 into the stable release channel, Google announced that it will now start enabling the relevance and measurement APIs in its browser. This will be a gradual rollout, with Google aiming for a 99% availability by mid-August." And here we are on the 15th, and my Chrome got it.

"At this point, Google doesn't expect," they said, "to make any major changes to the APIs. This includes virtually all of the core Privacy Sandbox features, including Topics, Protected Audience, Attribution Reporting, Private Aggregation, Shared Storage, and Fenced Frames. It's worth noting that, for the time being, Privacy Sandbox will run in parallel with third-party cookies in the browser. It won't be until early 2024 that Google will deprecate third-party cookies for 1% of Chrome's users." That's, you know, to make sure nothing unforeseen and horrible breaks. "After that," they wrote, "the process will speed up, and Google will deprecate these cookies for all users by the second half of 2024.

"The AdTech industry," they said, "has been able to test its readiness for the eventual third-party cookie deprecation, in part through the Relevance and Measurement origin trial. With these features moving into general availability, Google will end this trial, and revoke the tokens to run experiments on September 20th, 2023. For Chrome users, Google will now start rolling out its user interface" - and that's what has happened - "to allow them to manage Privacy Sandbox data in the browser, including ad topics, site-suggested ads, and ad measurement data. This rollout will run in parallel with the API releases. Google will soon make enrollment and attestation a mandatory process for

AdTech companies that want to access these APIs on Chrome and Android, though they will be able to continue to do some local testing, as well."

And finally they said: "Google notes in their recent announcement: 'Shipping these APIs is another key milestone in the ongoing Privacy Sandbox timeline. This marks the beginning of the transition from sites testing in the origin trial to integrating these APIs in production. We'll be keeping you updated as we progress through enabling the APIs, to the opt-in testing with labels in fourth quarter 2023, the 1% third-party cookie deprecation in first quarter 2024, heading towards the full third-party cookie phaseout in the third quarter of 2024.'"

Okay. So here's how Topics works. The essence of Topics are individual topic tokens - zero, one, or many - which are assigned to individual websites. For example, my GRC.com site might be associated with Computers and Electronics/Network Security, and Computers and Electronics/Programming, and Networking/Internet Security. So when someone visited GRC.com, their own web browser would record their interest in the topics associated with GRC.com, those topics, those three. But their visit to GRC.com itself would never be recorded other than in their regular local browser history as is always done. The only thing retained by the browser to indicate their interest in those topics would be those three numbered parameters.

For example, in Google's current 349-topic list, which they refer to as a "taxonomy," there's "Arts and Entertainment" as a general topic if nothing more specific is available. But then there's "Arts and Entertainment," and then under that "Acting and Theater," and "Comics," "Concerts and Music Festivals," "Dance," "Entertainment Industry," "Humor." And under "Humor" is the subtopic "Live Comedy." And it goes on like that with "Arts and Entertainment" having a total of 56 token entries before we switch to "Autos and Vehicles," which has 29 subcategories, which brings us to "Beauty and Fitness" and so on. You get the idea.

So here's how Google's specification explains this. They said: "The topics are selected from an advertising taxonomy. The initial taxonomy proposed for experimentation will include somewhere between a few hundred and a few thousand topics." They said: "Our initial design includes around 350." And I counted them, it's 349. "As a point of reference, the IAB Audience Taxonomy contains around 1,500 individual topics and will attempt to exclude sensitive topics." And they said: "We're planning to engage with external partners to help define this. The eventual goal is for the taxonomy to be sourced from an external party that incorporates feedback and ideas from across the industry."

Okay. Under today's tracking technology, if someone were to visit a website, for example, concerning the termination of pregnancy, that visit would be tracked and recorded by unknown and unknowable third parties. But under the forthcoming Topics system, the topic associated with that site might be "Health and Wellness" and nothing more specific. So this really is privacy centric. And as we'll see, Google has bent over backwards to ensure that the topics a user's browser volunteers cannot themselves be used as a tracking beacon.

Google explains: "The topics will be inferred by the browser. The browser will leverage a classifier model to map site hostnames to topics. The classifier weights will be public, perhaps built by an external partner, and will improve over time. It may make sense for sites to provide their own topics via meta tags, headers, or JavaScript, but that remains an open discussion for later." I have a link in the show notes to the official IAB Taxonomy in the form of an Excel spreadsheet, for anyone who's interested, so you can just sort of see what an existing advertising taxonomy looks like.

Okay. So as a user roams around the web, their browser infers the topics of the various sites they visit using this classifier which is built into the browser, which maps the site's

domain name to the topics that are relevant for it. For any given domain, the classifier may return nothing, no topics, one topic, or more. There's no set limit, though between one and three is what's expected. So it is possible that a site might map to no topics and so doesn't add to the user's accumulating topic history. Or it's possible that a site adds several topics or increases the "popularity weighting" of the user's existing topics by being another instance of their interest in health and wellness, for example.

Web browsers which support the Topics API divide the flow of days into discrete epochs which are each one week long. But it would be more precise to say "one week worth of seconds long" since there is no alignment to any calendar. Each browser instance chooses for itself randomly when each week-long epoch begins and ends. All of a user's browsing activity is grouped into these week-long epochs, and only the most recently finished previous three epochs are retained.

In other words, as soon as the currently open epoch closes at the end of its weekly cycle, or its week-long cycle, the Topics it contains becomes the most recent and is then available along with the two next most recent epochs, and the one that had been the oldest of the previous three is completely discarded. This has the effect of causing every user's browser to completely forget everything it had acquired and knew about its user on a rolling basis every four weeks. This is neat, since when a user's interests change, the Topics which reflect their new interests will realign with them automatically.

In the past we've talked about the DOM, that standardized Document Object Model, which is what our browsers, the way our browsers represent pages that we visit. One of the objects that's always present in this model, for example, is the "document" object. And it has properties like the "body," which is the document's body text; "images," which is a collection of all the images present in the document, and "links" which is a collection of all of the links present in the document. The Topics API adds the "browsingTopics" property to every page's document object.

When a website or one of its advertisers queries the document's "browsingTopics," they will receive up to three topics from the topic taxonomy, one from each - now, here, one from each of the preceding three weeks - and those three topics will be returned in random order. When a visitor visits a site, the site will be able to obtain up to three topics which might help it or its advertisers to choose a more relevant ad. It was decided to provide three topics, that is, you know, three, so that a site that doesn't see visitors often - that is, a specific visitor - will still obtain sufficient information about the visitor it doesn't see often to choose something useful. And a granularity of one week was used between a user's topic updates so that sites and advertisers which are seen much more often by a user's browser will learn at most one new topic per week.

And not all websites receive the same weekly topic from any given user. Here's how that works. For each week, once that epoch closes and the topics have been selected that were the most popular for the preceding week, for each week the user's top five, which were obtained from the web domains they visited during the preceding week, are determined using that topics classifier. And again, at no time does the user's browser contact any external servers for help with choosing. It is all done locally. It's all client-side. So five master topics are chosen from an examination of the preceding week's browsing history. And one additional topic is chosen completely at random from the entire taxonomy. So it will have nothing to do with the user's usage history at any time. It's just it's completely random.

Okay. Then, when the "document.browsingTopics()" API is called by a site, or typically an advertiser on a site that the user is visiting, or JavaScript running on its ads, the topic returned for each of the three weeks, the one topic returned for each one of the three weeks, which remember are also returned in random order, so that means nothing, is chosen from one of those top five plus that one random topic as follows, with a 5%

chance, so one in 20, the randomly chosen topic will be returned. And that's put in there as a deliberate wildcard.

Otherwise, the other 19 out of 20 times, the value of an HMAC hash is computed from a static per_user_private_key, so every single browser instance is unique; the week number, that is, as the browser is creating and closing these epochs, it will be incrementing the week number. So a static user private key, the week number, and the document page's website domain name. They're all mashed and hashed together, and the result is then taken modulo 5 to produce a uniformly distributed value from 0 to 4. That is used to choose one of the user's top five topics for that week. The same thing is done for the top five topics during each of the earlier two weeks. And it's those three topics which are returned in random order.

Now, this may seem overly complicated, but it's quite clever and privacy enforcing. The use of an HMAC keyed by a per-user secret, the week number, and the domain of the webpage means that, for that week, the user's browser will always present the same one-of-five topics, except for that 5% chance of the wildcard, to anyone querying from the same site; but that every different site will also see an unpredictable but constant for them one-of-five topics for that user for that week. This is done to minimize the amount of information being disclosed since this guarantees that no site will receive more than one fixed topic per user per week. And each site only ever receives one of the five real topics, which makes it impractical to cross-correlate the same user over time.

This 5% noise is introduced to ensure that each topic has a minimum number of members, as well as to provide some amount of plausible deniability, meaning that no topic can be regarded as absolute. It may have been deliberately chosen at random. And remember that the exact point in time where a user's week ends and the next one begins, that's fixed also, but chosen at random by their browser. So this introduces some additional uncertainty and noise, since not everyone's web browser will calculate new topics at the same time on the same day, nor will it be changing them.

Now, there's one last extremely tricky bit that's a bit of a mind-bender. But it's an important privacy safeguard for the entire system. It addresses the problem that FLoC had and which you, Leo, mentioned several times on other TWiT podcasts, and probably this one. The problem was that FLoC was broadcasting a token which was a condensation of the person's web browsing history, and thus by extension a condensation of them. And those who knew how to interpret the token would know what it meant. And this token was being presented to any website they visited without prompting. So you quite correctly identified this as introducing a significant reduction in user privacy. You'd go to a site you'd never visited before, and immediately your browser would be telling them a lot about you, even though you'd never been there before.

So this problem did not fall upon deaf ears, and the Topics API incorporates a mitigation for this. It's worth reminding everyone that there's nothing whatsoever salacious or even really very interesting about the list of topics. They are really quite bland and dry. But they make sense from an advertiser's standpoint. So the first point is there's just no way for anything very personal to be revealed or represented by these topics. But even given that, the Topics API incorporates a very strong filter. And here's the mind-bending part. I'll explain it, then I'll provide an example because the explanation won't do it.

Not every website that calls this API will receive all of a user's three chosen browser topics. Only API callers that observed the user's browser visiting some site which mapped to the topic in question that would be returned within the prior three weeks qualifies to receive the topic. I know, it's hard. In other words, if an advertiser on a website did not call the API sometime in the past three weeks for that user's browser, when they were visiting some site which mapped to a topic that would be returned now, then the topic will be filtered out and will not be included in the three-topic the array returned by the

API. Fewer topics will be returned. Since, as I said, this is somewhat mind-bending, here's an example.

During the previous couple of weeks a user has been browsing a lot about travel. So for the time being the browser has learned about them and chosen to represent their travel-related interest to the world as they visit other sites. So now suppose that they're at a site about gaming, and an advertiser on that site runs JavaScript in an ad insertion frame which queries the document.browsingTopics() API to receive the three chosen topics of interest to the user, and they of course come in no particular order, from among the topics that the user's browser has chosen to offer anyone querying about its user while they're on this gaming site. Remember that each week the top five topics of interest are chosen, and one of those five is selected from each of the previous three weeks based upon a hash of the domain name being visited.

Okay. So only if that advertiser, advertising on the gaming site, had queried that user's browsingTopics API sometime during the previous three weeks while that browser was at some other site whose topics matched the topic the user's browser had chosen to offer at this site now, would that topic be allowed to be returned and thus be presented to that advertiser.

In other words, in order to obtain an interest topic from a user when they are wandering around anywhere on the Internet, an advertiser must have previously asked their browser about them during the past three weeks while they were on a site whose topics may have contributed to the topic they are offering this week. This prevents the problem that FLoC had of recklessly blabbing about a user's interests.

Here's another way of thinking about it. In a world with Topics and nothing else, assuming that we've blocked third-party cookies, fingerprinting, and all other tracking mechanisms, or maybe outlawed it, the advertiser doesn't know who the user is. But they will have recently queried the user's browser while the user was at a website whose topics matched the topic the user is now offering. Google explains that this extra topic information filtering is intended to "prevent the direct dissemination of user information to more parties than the technology the API is replacing," in other words, third-party cookies and other tracking mechanisms. Another way of phrasing it is that this ensures that third parties don't learn more about a user's past than they could have with cookies.

Okay. As I've mentioned, the history window which limits the inclusion of topics available to each caller is three weeks. Only the topics of sites that use the API, or host ads that query the API, will contribute to the weekly calculation of topics. Moreover, only sites that were navigated to via a deliberate user gesture are included, as opposed to a redirect, for example. So it's not possible to bounce users around to load up their browser with topics. That won't work. If the API cannot be used for any reason, if it's disabled by the user or by a response header, then the page visit will not contribute to the weekly calculation. So it's easy for sites to opt out if they choose.

If the user opts out of the Topics API themselves by disabling it, which is trivial, it's a switch you can throw in Chrome. Or if they're in incognito mode, which automatically completely disables the Topics system, or the user has cleared their browser history, no topics will be returned.

So the goal of the Topics APIs is to take a step toward increasing user privacy, compared to our current state, which would not be difficult, while still providing enough relevant information to advertisers that websites can continue to thrive, but without the need for invasive tracking enabled via existing tracking methods.

One other crucial aspect of the whole Topics solution is the user's understanding and sense of control. Very, very few users are ever going to understand what we've just

described. They don't really need to. But what they will see is the list of the topics that their own browser has accumulated over the past three to four weeks. It's displayed right there in their browser, on that page. And they are completely free to either delete any topic they choose, or even permanently disable it from ever coming back.

Google wants this to succeed, so they want to give users of this system all of the control that they can imagine. The API's human-readable taxonomy, which is what is displayed on the page where you can just see the topics for yourself, enables people to learn about and control the topics that may be suggested about them by their browser. And as I said, they're free to delete or disable any of them. Clearly, one of the things that we all complained about with FLoC was that it was this bizarre, you know, completely opaque token blob that meant nothing to anyone unless you had the magic formula for what all the bits meant. And Google completely fixed that.

And also unlike FLoC, which built its hash from every site visited, only sites that include code which calls the Topics API are included in the browsing history eligible for topic frequency calculations. In other words, sites are not eligible for contributing to topic frequency calculations unless the site or an embedded service, you know, an ad or whatever, has taken action of calling the API.

So that's the operation of Google's Topics API, which is now in place in Chrome today, and presumably in other Chromium browsers which choose to implement it. It'll be there. They'd have to turn it off and remove it if they didn't want it. We know what the EFF would say. They want nothing less, they're satisfied with nothing less than pure and absolute anonymity. But website operators say that would cost them dearly. I have no way to judge how much revenue websites would lose if targeted advertising were eliminated. And I certainly wouldn't shed a tear over the end of any companies whose entire business model was secretly tracking users against their wishes and selling this information, you know, under the table.

To me, the Topics API feels as if Google is finally getting really serious about offering an alternative to tracking; one which, apparently, advertisers can live with and which will provide the websites which use it value. After all, Google is a massive advertiser themselves, and they plan to shut down Chrome's third-party cookies starting in 2024, with full elimination by the middle of next year. And if they will shut down the redirects that the use of Google Shirts creates, I would be very happy with that, too, because that annoys the crap out of me.

Anyway, once they've done that, we can imagine that Chrome will become extremely tracking hostile. Topics gives users not only a sense of control, but actual control, by allowing them to view, delete, disable the topics that are currently being judiciously doled out about them. And this brings us to an interesting question: If our personal interests are separated from tracking, so that the things we're interested in can be shared cleanly, and if that minimal sharing provides significantly greater advertising revenue to the sites we frequent, are we willing to voluntarily give sites that benefit? I know with absolute clarity that I am. If it's really that valuable, please accept it as my micropayment. You know, the EFF's absolutists will disagree with me. And that's fine. Topics finally makes that an individual personal choice that each user of Chrome can make for themselves.

Leo: So is the chief difference between this and FLoC the fact that it doesn't broadcast it? It's kind of private? Otherwise, it seems similar. It's like a modified FLoC to me.

Steve: So maybe anything that creates information could be considered a modified FLoC. What's different is there's nothing about it that's opaque, so it doesn't use this weird hash of websites that no one can interpret.

Leo: Right, right.

Steve: That was a big problem. So what Google has done is they have created a very simple - basically it's simple. The browser itself, Chrome, will know what few topics TWiT.tv, the domain TWiT.tv, stands for. Out of a list of about 1,500, there will be some number of topics, zero through three, that are associated with TWiT.tv. And that is built into Chrome.

Leo: You said 1,500. You mean 349?

Steve: No, 349 is just their current working R&D...

Leo: Oh, but it could be as much as 1,500. Okay.

Steve: Yes. Yeah, there really isn't a limit on the number, but the IAB currently operates with 1,500.

Leo: Yeah, they have 1,500, yeah. But it's ridiculous. It's demographics and stuff like that, that probably the website isn't going to show.

Steve: There will be no race, no gender, nothing.

Leo: So that eliminates a number of the IAB topics.

Steve: And so the point is that a - so Chrome knows what TWiT.tv is.

Leo: Right.

Steve: And you'll have from that taxonomy, you know, up to three things that are associated with TWiT. So when someone visits the TWiT site, the browser sees that their user has visited TWiT.tv and internally knows those three topics. This is this classifier. It is built into Chrome. All the domains are built into Chrome. And what it is that they're associated - and the up to three topics that they're associated with. So the browser dumps those topics into a bin. It just accumulates them. And at the end of a week, all of the topics associated with all of the sites the visitor went to, at the end of a week this epoch ends.

So at that point the five topics that were hit the most are chosen as the things the visitor cared most about in the previous week. And the five topics that were chosen from the week before are moved down to the second week, and the topics that were from the second week are moved to the third week. So we have three weeks. Each week has five

topics. So now when an advertiser queries the browser at a site, based on the domain where the advertiser is querying, the browser itself will select one from each of the five topics for each of the three weeks, deliberately scramble them, and say here are some things that the user of this browser is interested in.

And that's it. That's the entire system. Coupled with a user interface where the user is able to go and look at what those topics are and say, no, I'm not interested in that, and delete it. Or, oh, I don't ever want to be associated with this, and so disable it so that it can never be affiliated with them. So there's no tracking. Google has gone to extremes to disclose a minimal amount of still useful information about the user. And that's it.

Leo: Yeah.

Steve: And it's in there now.

Leo: This seems fair. I mean, honestly...

Steve: I think it is really - I think it is really fair, Leo.

Leo: This is how we sell ads on TWiT is, you know, your topic is tech. And in the case of your show it's security. And it's enterprise. And those might be the three topics.

Steve: Exactly, exactly.

Leo: And I have no problem with that. That's stuff that a human could deduce anyway.

Steve: Exactly.

Leo: So it's not - yeah. And it just automates something that a human would do anyway, I think.

Steve: Yes. And I think that the reason we're having trouble with this is that we have had such an adversarial relationship with, you know, the big bad Internet and how it wants, you know, and cookies planted and, you know, Panopticlick and fingerprinting and all of this crap. I mean, so it's difficult to say, wait a minute, you mean that's what this is? And Google's, like, that's what Google wants? It's like, yes.

Leo: There's a lot to be said for a privacy-forward system that still allows advertising to be topic-based.

Steve: Yes.

Leo: As you said at the beginning, who knows if that works or not? But at least we know advertisers demand it. So we're going to have to give it to them.

Steve: And that means that they pay sites for it.

Leo: Right. Yeah, I mean, we are advertising-supported. And we don't, you know, we don't have to worry about this kind of stuff because we're a podcast, not a website. But if you're The Verge, this is good news. This is something you want.

Steve: Yes, yes. And really for me I do have an interest in health and wellness. And so if saying yeah, I have an interest in health and wellness, if that's a micropayment that I can make to sites that I visit, I want to support them that way.

Leo: To be clear, you say that only by your behavior, by visiting that site.

Steve: Correct.

Leo: That's where it gets that information from. It's really what it's saying is it's categorizing sites. And then when you come to visit, now they've got - and are they attaching that to an identifier? They are; right? So it does follow you around. You visited a health and wellness site. So when you go to visit something else, it knows that's you. Yes?

Steve: Well, so you visit something else in that browser. So it's your browser...

Leo: Right. It's in the browser. Yeah, yeah.

Steve: It's all TNO. It's all local.

Leo: Right, right.

Steve: And that's the other cool thing about it is that it is transparent, and it is local.

Leo: Yeah. I think this is a good thing. I hope, I mean...

Steve: It is. I mean, it's like it's the answer we want.

Leo: And you're right, I mean, groups like EFF, which I support, probably come from the point of view that no ad is a good ad. And as a result, nothing's going to make them happy except, you know...

Steve: Yeah, who supports them?

Leo: Well, yeah, right. No, they're not ad-supported. I support them by giving them money. Which comes from ads on your show. So I guess, yeah...

Steve: And I support Wikipedia.

Leo: Same thing.

Steve: I get my little monthly notice, and thank you.

Leo: Yeah, that's what I do with EFF. Yes. This is good. This is very good. I'll be curious to see. I think at some point EFF and others are going to have to say, look, we're going to have something. Let's choose something that does the least harm. And this seems to be it. Finally.

Steve: Yeah.

Leo: This was Google's kind of plan all along was try some stuff, see what people said.

Steve: Yup.

Leo: That's why it was always in beta. Interesting. Now, you said it's going to propagate to all other Chromium-based browsers. This may not because it's not going to Chromium, I presume. Or is it?

Steve: I think it's in Chromium.

Leo: Oh, that's interesting.

Steve: I don't know. I think it's in Chromium. And now you can see why I want it in Firefox. It's not a bad thing.

Leo: No, no.

Steve: It's a good thing.

Leo: Yeah. They finally got to a point, it was almost a silent negotiation, where they finally said, well, okay. Will this work? No? All right. Well, how about this?

Steve: I think that's, yeah, I think it's exactly right is that they came up with something that provides a minimum of information, yet something useful to advertisers.

Leo: Well, we'll definitely be talking about it tomorrow on This Week in Google.

Steve: Yup.

Leo: And I will be looking more into it myself. Very interesting. As always, you know. And, boy, thank you for the very good news at the top of this show that you might - and by the way, you know, you're an at-will podcaster. If at some point, Episode 1004 you go, well, that was a terrible idea, I quit, nobody's going to stop you. Right?

Steve: Well, and as I said, I've been feeling this for a while. I just - it's hard for me to imagine, like, you know, killing something that's good. Like, you know...

Leo: Yeah. Thank you. Well, and needed. And you're doing something that's making a difference.

Steve: Everyone says so.

Leo: Yeah.

Steve: So I'm going to thank everybody.

Leo: Thank you for doing it. By the way, yes, it is in Chromium. I'm just looking at the Chromium site.

Steve: Cool.

Leo: It is part of Chromium. And we will definitely talk about it. And I think the world will be talking about the Steve Gibson promise that at least we're going to get an Episode 1000. I can't promise any more than that. Steve, you're the greatest. I appreciate it. Catch Steve at GRC.com. Just like on Episode 1.

Steve: Bye.