SECURITY NOW!

Transcript of Episode #755

Apple's Cert Surprise

Description: This week we reexamine the Windows 10 lost profiles problem, and also a consequence of the need to roll back (or avoid in the first place) the Patch Tuesday disaster. We look at a new feature to arrive with the next Windows 10 feature release, unfortunately named the 2004 release. We also examine the details of a new attack on the 4G LTE and 5G cellular technology, the full default rollout of Firefox's support for DoH, and also the availability of a powerful new sandboxing technology for Firefox. We also check in with Chrome's fix earlier today of a zero-day that was found being exploited in the wild. And, finally, before turning our attention to the bomb that Apple dropped in the lap of the entire certificate industry last week, I'm going to update our listeners about the things I've learned after returning to the work on SpinRite's next iteration.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-755.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-755-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with some fun Pictures of the Week, an example of how not to graph. Once again, news that makes me wonder why anybody's still using Windows. And he's going to explain why Apple suddenly and unilaterally announced that they're not going to accept certificates with a longer date than one year in Safari. Details, why it's happening and why it didn't have to, coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 755, recorded Tuesday, February 25th, 2020: Apple's Cert Surprise.

It's time for Security Now!, the show where we cover your security and privacy and all that jazz, and sometimes how things work, with Steve Gibson, our man about town, the man about the security certificate. He is the guy in charge at the GRC, the Gibson Research Corporation, at GRC.com and a beloved character around here because he keeps us up to date on what's going on in security. Hi, Steve.

Steve Gibson: Yo, Leo. Great to be with you again for this last podcast of February. Where has a month gone?

Leo: Well, it is the shortest month, after all.

Steve: Oh, actually I'm just looking at my little calendar. It says the 29th. So this is a...

Leo: Yes, it's a leap year. Happy Leap Year!

Steve: Yeah, 366 days this year. Which sort of ties in, in an odd way, to the topic of this podcast. I named it Apple's Cert Surprise because last week at the CAB Forum, the CA Browser Forum, thus CAB, Certificate Authority/Browser Forum, where browser makers and the certificate authorities meet regularly in order to talk about the world and how things are going and what things need to change and so forth. Apple did something that I have to wonder how truly unexpected it was because I imagine there were some people who were not as surprised as others. But I'll save that for the end of the show. But anyway, it bears on the number of days in the year because of the expiration of certificates, which is based on how many days in the future from the date of issuance the certificate self-expires. And of course our listeners know that certificate expiration and revocation and all that is like one of this podcast's favorite hobbyhorses because it's just so badly broken. But we have a lot of news to catch up on.

We're going to reexamine the Windows 10 lost profiles problem. It turns out there's news there. And also a consequence of the need to roll back, or maybe to avoid in the first place, this month's Patch Tuesday disaster. There are some consequences of that that need to be handled. We look at a new feature which is slated to arrive in the next Windows 10 feature release, which is unfortunately named the 2004 release, which, you know, okay, Microsoft, come on. Then we also examine the details of a new attack on the 4G LTE and 5G cellular technologies, which is worrisome. We've got the full default rollout of Firefox's support for DoH, and also the availability of a powerful new sandboxing technology arriving from Firefox.

We also check in with Chrome's fix earlier just this morning of a zero-day that was found being exploited in the wild, only the third of those to happen for Chrome in the last 12 months. And, finally, before turning our attention to the bomb that Apple dropped in the lap of the entire certificate industry last week, I'm going to update our listeners about the things I've learned after returning to work on SpinRite's next iteration. So I think we have another interesting podcast full of fun stuff.

And we have, for the Picture of the Week, this is completely off topic, but I labeled this our "Cranky Old Guy Pictures of the Week" because I saw something on Sunday that just so annoys me. I mean, like I stopped it and backed it up and pointed it out to Lorrie and said, "Look at this, how wrong this is."

Leo: Look at this. Look at this. This is wrong.

Steve: And I thought, okay, I've got to share this with our listeners. So, yeah.

Leo: I love it. All right. We're ready for the pictures.

Steve: So anyway, this is just - you and I have talked about in the past that one of my pet peeves is a chart where the Y axis, the vertical, I guess it's deliberately deceptive. Typically they're trying to show how great some change was over time. And you look at it, at this line going across, and you think, oh, my god. And then you look down, and you realize that the axis has been labeled so that, you know, it doesn't start at zero, so that you're seeing an actual percentage of change over time. Instead, the minimum of the Y axis is just a little bit lower than the maximum, which has the effect of ridiculously inflating the size of the change.

Well, in a related peeve, I was watching "Meet the Press" on NBC on Sunday, and they were showing what was trying to be a pie chart of an entrance poll from the Democratic Nevada Caucus the previous day. And what just stopped me cold was that the first of these pictures that I show, show that Sanders and Biden both have the same percentage in this entrance poll.

Leo: Oh, really.

Steve: Uh-huh. They're both 24%. And you look at this chart. And so the problem is it's technically a pie chart, but now I guess it's more fancy to do it as a doughnut where you remove the center; right? So but the whole point is that our eyes are very good at discerning angle. This is actually one of the reasons that speedometers use an angular needle. And even in high-performance, high-end attack military aircraft, they still have dials. They have pointers because we're able to instantly understand an angle, whereas otherwise we'd have to be, like, reading digits and thinking about it.

Well, so the beauty of a pie chart is that it's all about angles. But in this case, and we see it all the time, in their interest to make it look fancy and modern, they, of course, they make it 3D, and then they drop it back. They tilt it backwards so that it looks more 3D-ish and fancy. And of course it has the effect of, because of perspective, of dramatically increasing the foreground versus the background.

Leo: It's kind of like a Mercator projection or something. I mean, we all have this...

Steve: Yes, that's a very good analogy, yes.

Leo: It's the opposite problem, which is taking a 3D space and making it 2D misrepresents the size of the continents and the countries. In this case, taking a 2D space and making it 3D misrepresents it. You know, conspiracy theorists are going to say, well, mainstream media doesn't like Bernie anyway, so they're definitely trying to make Bernie look smaller than Joe. And it's true, they could have put Biden in the back and Bernie in the front.

Steve: Well, for what it's worth, I was going to say it's not even alphabetical. But for those who are not able to see this, the first chart shows Sanders and Biden, the numerical percentage is 24 each.

Leo: Yeah.

Steve: But Biden's slice of the pie looks...

Leo: It's huge.

Steve: ...dramatically larger than Sanders's. And so I first thought that. I go, oh, my god. And I backed it up and, you know, subjected Lorrie to the same tirade. And then a few minutes later we get the second one, which is also a disaster. Here, again because of this distortion, we've got another Sanders/Biden comparison. And in this instance, their size looks sort of comparable. It looks like, okay, that's kind of about the same. But in fact Sanders is at 27, and Biden is at 39.

Leo: The numbers are not at all comparable.

Steve: Not at all comparable. Yet, you know, if you just sort of look at the amount of the color...

Leo: The blue should be almost 50% bigger than the purple.

Steve: Yeah.

Leo: But it's not.

Steve: No.

Leo: No.

Steve: Anyway, so...

Leo: So don't believe these charts. Look at the numbers. Make your own charts, kids.

Steve: Yeah. Or, again, if it was just a straight-on pie chart, we would instantly have a deep intuitive sense for the relative size of these things. And that's the whole point of having an angular pie chart. It's to show you relative sizes. But it's completely destroyed when you throw it into 3D and then get perspective distortion in order to make it look really fancy. So, bad. Bad NBC.

Leo: Bad. Fake news.

Steve: And speaking of bad, so remember how last week we said that even though after installing the February - this is the infamous now KB4532693 update rollup. Some users discovered that all their stuff was gone. And their desktop, all the icons disappeared from their desktop. They went to My Documents, and there were no documents there under My Documents. Everything was gone.

And we know that this is the per-user profile information, which is when you log onto Windows as a user, that's what Windows makes current. So you're seeing under the user's directory off of the root of your boot drive is your various usernames who share the computer. You log in as a user, you get your stuff. Somebody else logs in, they get their stuff. Unfortunately, you run the February update rollup, and in some cases you get no stuff.

Well, the good news, we believed, was that nothing was actually deleted, but rather the user who did the update had their profile temporarily renamed with a .000 or a .bak, which meant that the stuff wasn't gone, it could be recovered. Some advice was you could uninstall this. Oddly enough, apparently, if you just restarted Windows four to six times - and I can't believe I'm saying this because it just seems so broken. But yes, children, you just restart Windows, just keep at it up to six times, and maybe all your stuff will come back because something in Windows will wake up and go, oh, whoops, and then bring it back. Turns out that's not true for all people.

What is beginning to surface, and there are enough people who seem authoritative enough, not just confused newbies, but people who appear to know what they're doing, reporting that in some instances they have actually permanently lost everything. That is, Windows did rename their profile into oblivion.

There's a posting on the Answers forum at Microsoft.com that reads: "I was on the phone with Microsoft for four hours. Surface Pro 7, about two months old. There's no user data, no temporary account, no restore points. Uninstalling the update didn't work because my user data is gone. Microsoft is calling in the morning because they think my personal services are free." So this guy's a little peeved, obviously. "I know the only option is to flatten the Surface and start over, like 'Groundhog Day' the movie. I have two desktops upgraded from Windows 7 that survived, if that helps."

Somebody else also posting in the forum: "This update KB4532693 caused all the data on my laptop to be erased. Even after uninstalling the update, the laptop would not successfully boot. Then resetting the laptop to factory settings while choosing to keep all personal data erased everything. EVERYTHING," he repeats in all caps.

And there are major news reporting groups like Bleeping Computer that have followed all these threads down, and they're confirming that, as far as they know, this is true. Bleeping Computer's article has the headline: "Windows 10 KB4532693 Update Bug Reportedly Deletes User Files." And again, we're not really getting anything from Microsoft. And at this point it would almost be - it would be helpful if Microsoft would say something about this because this becomes a problem.

As we know, it's really the case nowadays that our backups need to have backups. And after I had a few close calls with old and spontaneously dying machines - remember that my beloved Windows XP machine died, just completely died, and all of my stuff was on a hardware RAID with multiple drives, I mean, it was all there still, and I got it all back, and everything was fine. But still it was like, yikes. I mean, it sort of was a wakeup call for me. So now I am backed up every which way. And in fact I'm seriously considering moving my entire Windows User Profile now to Sync.com, rather than just selected working directory hierarchies, which is what I've done. I mean, I just continue to be so pleased with the way Sync.com is working. And that way all of my user profile stuff would be synced and version managed across multiple machines and so forth.

You know, I firmly practice safe computing. I'm careful about where I go. I don't venture very far away from places on the Internet that seem really safe. And I'm very careful when I'm, like, needing to go download something. I spend some time checking it out. I'm sure that our listeners do, too. So I hope that the danger from anything malicious successfully attacking me is minimal. But it's disturbing to imagine that running a Windows Update could be the Trojan horse that I willingly allow to enter through my front gate, and that it would do some damage. I mean, I've not lost anything. And we know that the majority of people aren't.

So what's unknown is, like, what exactly is going on? As long as, like, if it's the case that people are actually having all their stuff deleted by Windows Update, then that suggests that we really do need to perform, I mean, I'm thinking that I will do a full image of my system before I do an update, and that I will take control of that process rather than just let Windows do it whenever it wants to. So, wow. It's disturbing that this could be the case. And at this point, if Microsoft would say we know what's happening, we found that and figured it out, then that would be helpful.

And Leo, I'm unable to listen to all of Windows Weekly as often as I would wish to, but I did, while I was digging into all this, I encountered something that I remembered you guys talking about, I think, which was that Microsoft let go a huge percentage of their QA staff, and that now the developers are doing the testing of things that used to be a whole separate group. Is that the case? I think I remembered hearing that along the way. So anyway. Let's just hope that this gets resolved. And it would be nice if we heard from Microsoft, like some actual resolution to this issue.

Okay. So also, remember the serious zero-day exploit that we learned about which involved IE and its invocation of this older JScript.dll. We know that this vulnerability has been exploited in the wild in limited targeted attacks. The bad guys are able to leverage it to silently execute arbitrary commands on an unpatched system when the user visits a specially crafted website. And then we've subsequently learned that there are many other ways to get IE to invoke that old retired JScript.dll which is no longer the one that anyone's using, but it's still around because it turns out it's important to have it.

The issue is so severe that Microsoft was prompted to suggest a short-term fix until the February Patch Update became available, which would fix it for real. That flaw is being tracked as CVE-2020-0674. And the temporary fix was to remove all access to that DLL by deleting all security permissions from the file so then nothing, whether good or bad, could cause it to be invoked. It would just be off-limits in the file system. But as we know, many people soon discovered that their Windows Media Player, that was the thing that I think most people, and it's what I heard about most, would no longer work. But then other things were broken, too, sort of some obscure things like HP printing and other USB printers would no longer function after that DLL was taken out of service.

So the solution, we believed, was going to be Patch Tuesday's rollup, which would fix this for real. But of course now we know that some people have had to back out of that. I would imagine that enterprises would be terrified to apply this Tuesday rollup to their systems without extensive testing. So this of course leaves us vulnerable to a targeted attack which is probably going to be more prevalent now than it was because before the bad guys who were using it knew that it had been discovered, they were trying to keep it secret. As soon as it becomes clear that the window of opportunity of exploitation is going to close, then they can afford to be much less cautious and do as much damage as they can before this problem gets resolved.

Well, the problem is that its resolution was Patch Tuesday, which we no longer really have for many people who have had to back out of it or are, you could argue wisely, deferring its installation, maybe just skip it completely and hope that March turns out to have a better outcome.

Leo: Or use Linux.

Steve: Exactly. In fact, there were actually - I chose not to share them on the podcast because in some cases I could not repeat the language that was being used. But, I mean, there were many people who were so infuriated, I mean, like, people who clearly knew what they were doing, who were long-term Windows users. Anyway, the point is they were swearing off of Windows. They said, that's it, I'm not putting up with this "C" word or "S" word any longer. That's it. I am, you know...

Leo: It's a fairly serious flaw in an operating system to seem like they've lost your stuff. I mean, that's not insignificant.

Steve: Well, and the point was that to some people, Leo, and apparently Microsoft in interactions on the phone has had to admit that, yeah, uh-huh, it does look like it actually is all gone.

Leo: Ohhh.

Steve: That's what Bleeping Computer is now reporting.

Leo: Oh, no. That's it. I would not, you're right, I wouldn't use Windows ever again.

Steve: Yes. It is true...

Leo: You lose my data, you lost my business. That's ridiculous.

Steve: Yes. At this point. And so that's what's happening.

Leo: Oh, my goodness.

Steve: Yes, is that some users who apparently know what they're talking about and have spent hours on the phone with Microsoft have confirmed their stuff is actually gone.

Leo: Wow.

Steve: It wasn't just deleted. It was renamed into oblivion, and it was lost.

Leo: Oh, that's so not good.

Steve: Yeah. Yeah. So, okay, so in the case of this zero-day exploit, which we were hoping February's patch updates would fix, because we don't have those, the guys at 0patch.com, remember numeric "0" patch dot com, they do these micropatches. So I just wanted to let our listeners know that they have added support for 1903 and 1909 to the support for the fix for this JScript.dll. It is 0patch.com. I've got the link in the show notes for anyone who wants it. But I'm sure you can find it at 0patch, numeric "0" P-A-T-C-H dot com.

It wasn't initially working for the most recent two updates to Windows 10, the major feature updates. 1903 and 1909 weren't supported. It was Windows 7, Windows 10, 1709, 1803, and 1809. Now these latest two have been added because users may be relying upon it if they don't feel comfortable installing this February Patch Tuesday update. So anyway, that now exists.

Leo: And you trust these guys for their little - I mean, we've talked about this before. I just want to reiterate.

Steve: Yeah, I mean, their hearts seem to be in the right place. It's free for noncommercial use. Commercial users are asked to pay $25 per agent per year. And there is an enterprise plan. Yeah, I mean, I see no reason not to trust them. Often you don't have to reboot your system. They're, like, little 25-byte patches that just go in and fix this one problem until Microsoft comes along with an updated DLL that then fixes it for sure. But so here's the problem. We have this zero-day we know is being actively exploited in the wild. Yes, it has been targeted. But it's now known to have a limited shelf life because Windows is onto them, Microsoft is onto them. Windows is going to get patched.

It was supposed to be fixed in February. So for most of us it's fixed. And I'm now, like, I'm glad I didn't lose all my stuff when I did the February update. I'm going to be doing, you know, I've got good imaging tools all over the place. I'm going to do an image of my system from now on before I do a Patch Tuesday update because...

Leo: I guess that's sensible, but what a pain.

Steve: I know. Isn't that wrong?

Leo: Crazy.

Steve: And what's worse is that Microsoft isn't saying, oh, we know what happened, and we have the fix for it. And Leo, I think it was on Windows Weekly that you and Mary Jo and Paul were talking about how Microsoft had changed the way they're vetting these things. I think what I read while I was doing this research, and I think I recall you guys talking about it, was Microsoft laid off a huge percentage of its QA staff, and the vetting of these things was now going to be done by the developer group at Microsoft.

Leo: Yup. That was a while ago. And, boy, are they getting bit by that, I think.

Steve: Yeah, exactly. I mean, how many times have we talked about the fact that developers are unable to test their own stuff? You know, we can't test our own stuff. It's why I had this fabulous community in GRC's newsgroups, specifically to help me find my problems. Because you just can't find them yourself. You have to have other people who have different systems.

I dug into this a little bit because I was interested in how this was happening, because this is now a monthly problem for Windows 10. We're seeing now a run of problems with Windows 10 updates where they just seem unable to get it right. And apparently they're also relying on their insider community, who are not their developers, but they are users. But insiders are typically testing on VMs, and so they're not testing on mature, everyday use systems. And the problem is Microsoft is not saying, oh, we've found the problem and fixed it. So we're re-releasing this. No one need worry about it. Instead, there are hundreds of posts now of people who understand the temporary rename problem, and that's not what they have. They have permanent deletion of all of their data.

Leo: That's unacceptable. That's just - you can't get worse than that. That's terrible. Oh, my god, that's awful.

Steve: I know. Yeah. So, I mean, yeah. And imagine an enterprise environment. Again, until we know, it's one thing to say, oh, we figured it out, it turns out it's like there was that problem, that interaction with Adobe Creative Cloud stuff. It was the Adobe Creative Cloud services, blah blah blah. You know, great. Now we know. But this is just now like, as you said, I mean, obviously you understand. Who wouldn't? If a Windows update deletes all of your stuff?

Leo: It can't get any worse.

Steve: And most users no longer have a choice about whether to install the update or not; right? I mean, Microsoft has also said, well, we'll let you put it off.

Leo: Yeah.

Steve: Yeah, we'll let you put it off for a few weeks. But then we're going to force you to take it. And what if it forces you to lose all your data?

Leo: Wow, yeah. I think you have a court case. Criminently. Criminently.

Steve: Okay. So here's what I don't get. We've got the next major feature release for Windows 10 is the ill-named 2004 release.

Leo: I'm sorry, Steve, that's twenty oh four. Twenty oh four.

Steve: Oh, yeah, right, yes, right, twenty oh four.

Leo: Yeah. It is a little confusing because it kind of looks like 2004. Which seems like a pretty old update. But okay.

Steve: Well, yeah. I mean, if it's 1903, then you know that that doesn't refer to probably a year when you've been alive. So that's not confusing.

Leo: Well, honestly, everybody, in Windows Weekly we're just like, ugh. Everybody thought, oh, why don't they just call it 20H1, first half of 20, and then the second one 20H2? That was their internal code name. But no, they really wanted 2004. And by the way, it's not like it's that meaningful because they never make the date that they say. It's not going to be April. It's going to be whenever, you know, probably May or, who knows, March.

Steve: As long as they get it right. I mean, please. Hey, Microsoft, take your time.

Leo: Yeah, you're right, no hurry. No hurry.

Steve: Not on my account.

Leo: Nobody's asking for these. Nobody is asking for these, by the way.

Steve: I know. And here's another reason. They're now allowing optional device driver updates.

Leo: Oh, god, I hate Windows.

Steve: Starting now, as in like, what is the day?

Leo: What does that even mean?

Steve: I know. Get this, Leo. Device driver providers will be able to mark their device driver updates as "automatic" or "manual." Any device driver update marked "automatic" will be included in the Windows automatic update package, like it is now. But any device driver updates marked "manual" will first appear in the new Optional Updates section to be added to Windows 10. Now, okay.

Leo: See, my problem is, like if you're a business with an IT department, good. They'll handle this. But home users, stop buying Windows. Stop your friends; stop your relatives. I talk to way too many very nice people on the radio show who should never have run Windows. The word must go forth from this day into the future. Stop buying Windows for home use.

Steve: And, okay, so think about it. All of us who have used previous versions of Windows, like Windows 7, there is this odd - there are these two tabs; right? There's important updates and optional updates. And so you look at it, and you go, well, they're updates. Do I need them? And when you look at the optional ones, they're like security things. It's like, Windows Security Patch, blah blah blah. And it's like, oh. Well, that sounds like a good thing.

So, I mean, think about it, Leo. How can it be optional? I mean, like, what is it? And so it seemed to me to be a step forward that with Windows 10 it was going to - we're going to get a roll-up, and we just do it once a month, and it's going to be better. But they couldn't live with that. They couldn't stay with that. So they've, like, for whatever reason, decided to peel off, and now they're going to expose to the user this new, starting with the Windows 10 2004, which is beginning to seem more like 2004, you know, 16 years ago, this new optional update where, like, the user can select which they want.

But, okay, what should they do? I just don't get it. Microsoft explains that this change will allow hardware developers to roll out new drivers and test them for reliability - okay, shouldn't they have done that first? - against a smaller group of Windows users before pushing them out to a wider audience.

Leo: Well, that makes sense.

Steve: Oh, well, okay. They said Microsoft believes these changes will help their customers to "get the highest quality and most reliable drivers faster and with less friction." Okay. They're not going to back port this change to the earlier versions of Windows 10. So anybody with earlier version of Windows 10, you go to the driver manager app and, what, then click on Drivers and then go into the Properties and go to Update, and see if there's an update for that hardware, for the driver for that hardware. What? Like, who does that? It is just, well, yes, Leo, I think you clearly articulated the proper philosophy, which is Chromebook.

Leo: Normal people shouldn't use this, yeah.

Steve: Chromebook.

Leo: Yeah. And if you need a real operating system, Linux is very stable, reliable, easy to use, and has to date in my 30 years of using it never deleted my own directory, ever. Holy cow. Oh, geez, Louise. Okay.

Steve: Yeah.

Leo: Wow. What a world.

Steve: Yeah, Bleeping Computer reported that users were reporting...

Leo: That's terrifying.

Steve: ...permanent deletion of all their data.

Leo: I kind of understand the optional updates because they're basically giving the hardware manufacturer more control over whether you get updated. But there seems like there'd be other ways to do that than this complicated end-user dance.

Steve: But I guess I don't understand. It's either unnecessary or necessary. How is it...

Leo: Well, because they want to be able to offer optional updates to beta versions, to test it. They don't want to push - I don't know. You're right. This doesn't make any sense. I'm sorry. I'm trying to.

Steve: Yeah. And now they're going to give that to the end user. You're going to start getting calls.

Leo: I know. Is it optional? Should I do it?

Steve: On The Tech Guy, yeah. This says "optional updates." Well, how optional are they? Are they really optional, or not really optional? Because I'm not in a hurry today, and so I've got some extra time. And are they going to pile up? Are they just going to keep accumulating? Do they go away?

Leo: Well, that's a good question, yeah, yeah. Do you get another optional one? Or when does it become non-optional? At what point does that happen?

Steve: Right. You know, it's like user choice that users should not have to choose.

Leo: Right. It's just more BS. All right, Steve. On we go.

Steve: So we've got problems.

Leo: You mean we didn't before?

Steve: We've got your problems right here. The guys that have been poking at cellular phone security for years are today, the 25th of February 2020, during the NDSS, which is the Network Distributed System Security Symposium, which I guess should really be NDSSS, but perhaps they thought that that was one "S" too many. That's being held in San Diego right now. They're delivering their paper disclosing their new attack on 4G LTE; and, unfortunately, it also works against 5G because we haven't actually solved the problems.

Okay. So there is in our cellular system the same sort of hierarchy of levels forming a network stack as we're used to having, for example, in our computers with TCP/IP, physical layer, transport layer, and the various protocol layers and so forth. So even their abstract of their paper was just - it required too much understanding. I read it, and I thought, okay, well, I could tackle explaining this, but I don't think we have enough time. But the introduction gives us a good sense for the importance of what they have found.

So in the intro to their paper they explain: "Long Term Evolution" - which is of course what the acronym LTE or the abbreviation LTE stands for - "is the latest widely deployed mobile communication standard and is used by hundreds of millions of people worldwide. The protocol offers high-speed Internet access and packet-based telephony services and has become an integral component of our daily communication. We fundamentally rely on the security of LTE for a variety of applications. The security goals of LTE include, amongst others, mutual authentication, traffic confidentiality, and location privacy. Any attack vector undermining these security aims has far-reaching implications to the use of LTE as a communication medium.

"In the context of mobile communication, mutual authentication is an important security aim since it ensures that both communication parties, the user equipment and the network, mutually verify their identities." That is, you know, we're sure we're communicating to the party we think we are and vice versa. "As the wireless medium is accessible for everyone in the vicinity, and identifiers can be easily forged, mutual authentication is essential for building trust between communication parties. The telecommunication providers rely on user authentication for accounting, authorization, and the association of data sessions to a legal party. The latter case is of particular importance in prosecution" - and in fact this does, as we'll see, have some significant implications for defense attorneys - "in which a possible offender is accused of committing a crime via a mobile Internet connection. Additionally, users rely on network authentication for the confidentiality of their communication.

"One important example for missing network authentication is the second mobile network generation GSM (Global System for Mobile Communications). By faking the identity of a legitimate network, an attacker can impersonate the network in GSM and eavesdrop on the communication of the victim." And of course we well know that that's the famous phenomenon that we see in Las Vegas during the hacker conferences, where the number of apparent cell towers jumps by a factor of 10 overnight. And it's like, wait a minute.

So they said: "In contrast to earlier network generations like GSM, LTE establishes mutual authentication on layer three of the network stack using a provably secure Authentication and Key Agreement (AKA) protocol. Based on this protocol, subsequent encryption ensures the confidentiality of user and control data." They said: "Permanent integrity protection, however, is only" - permanent integrity protection, that's the key. On one hand we have encryption, which gives us privacy. But remember that integrity is the second side of that. We need authentication. We need something. We need to know that nothing has been changed, which is integrity protection.

They said: "Permanent integrity protection, however, is only applied to the control data. A recent study has revealed that missing integrity protection of the user plane on layer two allows the manipulation of user data in a deterministic way. Specifically, a layer two attacker in a man-in-the-middle position between the phone and the network can introduce undetectable bit flips due to malleable encryption and redirect traffic to another destination. While this attack demonstrates the potential consequences of traffic manipulation, it is solely limited to redirecting traffic to another destination."

So that was prior work upon which these guys based their next-gen attack. They said: "In this work, we introduce a novel cross-layer attack concept that complements the known two-layer vulnerability, that is, the missing integrity protection on the user plane, with exploiting the default IP stack behavior of operating systems on layer three." And specifically they target iOS and Android, both which are used in successful attacks. They said: "We make use of the reflection mechanism of certain IP packets," and I'll cheat and just say that that's ICMP Ping and ICMP destination Unreachable packets.

They said: "Exploiting the default behavior on these operating systems, we make use of the reflection mechanism of certain IP packets, which allows us to not only redirect user-plane traffic, but also to create an encryption and decryption oracle that enables an adversary to perform a full impersonation of the phone or the network on the user plane," meaning to appear as the user to the network and to appear as the network to the user, thanks to decryption. "We call this concept I-M-P-4-G-T," which is IMPersonation in 4G neTworks, pronounced "impact." "IMP4GT completely breaks the mutual authentication property for the user plane on layer three, as an attacker can send and receive arbitrary IP packets despite encryption."

And I'll finish just by wrapping up this part: "This attack," they say, "has far-reaching consequences for providers and users. Providers can no longer assume that an IP connection actually originates from the user. Billing mechanisms can be triggered by an adversary, causing the exhaustion of data limits, and any access control or the provider's firewall can be bypassed. A potential impersonation also has consequences for legal prosecution, as an attacker can establish arbitrary IP connections associated with the victim's identity."

Okay. So I mentioned briefly, they mentioned, this encryption-decryption oracle. That's the key to this. They establish a man-in-the-middle interception using a software-defined radio, which are now widely available. They are then enabled to probe the encryption by flipping bits, which results in a failure and a retransmission. So it is seen as an over-the-air problem, not an attack, which causes the endpoint to retransmit. They inject ICMP Unreachable and ICMP Ping packets into the stream in order to get either endpoint to reply. And since we've talked about, especially in the early days of this podcast, a lot about encryption, they explain the encryption and the decryption oracle operation, and it's understandable.

They said of the encryption oracle: "The goal of an encryption oracle is to learn the keystream of a connection, which later allows us to encrypt and inject arbitrary packets. For encrypting a target plaintext, the oracle injects a known plaintext into the system. The system encrypts the packet by XORing the known plaintext with a valid keystream for transmission, which is returned to the oracle."

Okay. So what they just said was, I mean, it's so obvious. They take a known plaintext, I mean, it could be all zeroes for all anyone cares. But it probably needs to be, for example, an ICMP Ping. But they know what it is. They inject that into the stream. The recipient encrypts it by XORing it with the keystream and returns it. Well, we know how to remove XORing; right? We re-XOR what we got back with the known plaintext, and that gives us the keystream, which is the output of a stream cipher, which the communications uses as an XOR pad, essentially, in order to create an encryption, which is solid as long as it's never reused. So this is a means of obtaining the keystream from the cipher in a way that then allows the attacker to reuse it.

They said: "Now, the oracle can extract the valid keystream by XORing the known plaintext on the encrypted packet. Any arbitrary payload can now be encrypted by XORing the target plaintext and the keystream," which has been determined. So basically that bypasses the problem of this data being encrypted on that layer. They said of the decryption oracle: "The goal of a decryption oracle is to decrypt and access the payload of an encrypted packet. To achieve the decryption of a packet, the oracle manipulates the to-be-decrypted ciphertext and sends it to the system. The system decrypts the packet and subsequently sends it back to the oracle. In this way, and similar to encryption, we can receive the plaintext of encrypted packets."

So they go into far greater detail in their paper. I've got a link to it in the show notes for anyone who's interested. It's being delivered today, as I mentioned, in San Diego. But they have conclusively demonstrated a fundamental weakness. We already had it in GSM. We knew it was a problem. We thought, oh, we fixed this problem in 4G LTE. Except we didn't. Nor is it fixed in the forthcoming 5G, since neither of these systems provides the needed message integrity protection at the user layer, which is where this exploit happens. It must have been assumed by non-cryptographer designers that the encryption running at the user layer would be sufficient to protect the user's communications. That is to say, they must have assumed that the encryption running in the user's application layer would be sufficient.

But it turns out there are available exploits. We know that XOR-based stream ciphers, while highly attractive due to their economy and the ease of implementation, are also highly susceptible to interception attacks that can trivially reveal the keystream if the plaintext can be known. You just XOR what you know the person talked about. And we've encountered and have talked about various attacks through the years on this podcast against simple XORing of stream ciphers. These guys clearly state that the only way for this to be fixed is for all of our existing cell system infrastructure hardware to be upgraded at the smartphone and the cell tower level. And we all know that's never going to happen. They are hoping that there might still be time to head off implementation of 5G, which repeats these mistakes. But they acknowledge it's unlikely.

So what it means to us is that application-level services like iMessage and Signal, and WhatsApp for that matter, which provide their own application level encryption and secure management, they're secure against this for, you know, their own end-to-end encryption. They don't rely upon the integrity of the underlying channel. But HTTPS is less certain because there we're relying on, just for standard web browsing, we are relying upon some aspects of the integrity of the underlying network. We are assuming that DNS is giving us the right IP, and that we're actually connecting to the machine at that IP that we think are, and that its certificate has not been spoofed. So we're trusting its certificate.

We know that certificates can be obtained, presumably by state-level actors, on a whim. They still have the problem of getting us to a spoofed server using a spoofed certificate. DNS and IP switching integrity is what we rely on. That's what this system is able to subvert. So this could be the component of a targeted attack by a state-level actor against specific individuals in specific settings. The good news is it's not deployable at scale over the Internet. This requires physical man-in-the-middle proximity. And it is a sophisticated attack, so it would only be targeted. And it needs multiple layers in order to redirect web traffic.

I'm pretty sure that without a lot more work the higher level encryption protocols that are providing end-to-end encryption through multiple strong endpoint encryption, in other words, iMessage and Signal and WhatsApp and so forth, you know, those guys, they're probably safe. But this is now a new attack, revealed today, and we're going to be using 4G LTE for a long time. Again, most of us have nothing to worry about. The use of strong messaging keeps us safe. But it would subject standard web communications to an attack by a state-level actor who is able to put all the pieces in place, and there are many of them. But they would be able to pull off site spoofing in a way that was undetectable.

Well, except by Chrome, actually, because I don't think you can spoof, well, maybe - have to think about whether you could spoof - I don't think you could spoof the serial number on the certificate. So it would still be limited, depending upon what sites were being visited. It would only be - Chrome could only detect it if you were going to Google properties where it knows what the certificate is, has been pinned for, Google. But still, you know, this suggests that we're unable to completely take the connections we have to our cellular providers as secure by default. So, yeah.

Starting today, also in today's news, we have Mozilla's rollout of DoH by default, turned on, in Firefox. Starting today, any new installations of Firefox will have DoH enabled by default. That's of course DNS over HTTPS. Then over the next few weeks it will be silently enabled for all Firefox users in the United States. They're going to gradually roll it out so that if there are any showstopper issues that are discovered that haven't already been found - basically it's been an overwhelming success so far. So no showstoppers have been identified.

In the U.S., over the next few weeks, the rest of Firefox users will be switched over to DoH. The only users who will not receive this update are those who have specifically disabled DoH in Firefox's Settings previously. And as we know, we've covered this on a number of podcasts, the move to encrypting DNS by tunneling it over HTTPS has not been welcomed by everyone, by a long shot. The most concerted pushback came from the U.K., where remember that ISP association went as far as to nominate Mozilla in 2019 as the year's Internet Villain, due to its work on the DoH protocol. And of course they regretted that nomination and subsequently rescinded it.

But the ISPs warned that rolling out DoH would cripple the U.K.'s national firewall system, which ISPs and law enforcement are using to limit access to child abuse websites and copyright infringement domains. After their lobbying efforts were joined by law enforcement and the British government, Mozilla themselves capitulated last summer, in July, and announced that they would not be enabling DoH for U.K. users, at least for now. I presume individual U.K. users can also enable it, if they choose. But it's just not going to happen by default.

So whether or not ISPs and governments like it, DoH appears to be where the industry is headed. So I think this is going to end up being a transient upheaval, and those who are against it are going to end up losing. It is now, as we last talked about this, supported by all major web browsers. Even if it's not always easy to find the setting, it's in there. And even Microsoft has announced plans to support DoH natively in Windows in the future. So whereas right now, if you're using a browser with DoH DNS resolution on Windows, only your browser is getting the advantage of the encryption and security benefits of DoH. If Microsoft moves it down into the OS and supports it natively, then all browsers, whether they're configured - as long as they're using the OS's DNS resolution, which is the typical default, and all other things, email and everything else on Windows, would get DoH. And that suggests that we're probably going to see this go OS-wide at some point.

So it is still Cloudflare the default provider of Mozilla's DoH. And that's a decision I wholeheartedly support. Not only has Cloudflare formally asserted that they will not log nor monitor nor intercept this use of DNS over HTTPS in any way, but it's Cloudflare who is making that assertion, and there are few companies that I would trust more to actually honor their pledge.

Leo: Good.

Steve: Yeah. And if somebody for some reason doesn't want to use Cloudflare, just wants to be contrarian perhaps, Mozilla does offer, Firefox offers the selection of NextDNS, which is also another good group. And so you can choose that if you'd rather not use Cloudflare. So anyway, you'll get it if you haven't deliberately turned it off. I've turned mine on, the moment we started talking about it, the moment it appeared in the UI. And my Firefox works just as well as my Chrome does. So I don't see any downside. And I'm not that worried about Cox, my ISP, my cable service, snooping on my web browsing, but they can't do it when I do it under Firefox. And that's going to be the case for everybody by default, any new installations starting today. And within a few weeks everybody should have it.

Leo: Cool.

Steve: And in another piece of important Firefox news, we've got a new generation sandbox coming first, interestingly, on Firefox for Linux and Mac, and only a little bit later to Windows. It's the result of a bunch of hard work by a team from UC San Diego, UT Austin, Stanford University, and Mozilla. It brings the next step in protecting users, both from malicious and inadvertently exploitable libraries that their web pages may load. There is a paper which the group has published. It's up on GitHub. I've got the link in the show notes. It's titled "Retrofitting Fine Grain Isolation in the Firefox Renderer." And to give our listeners a sense for what they've done, I'll share from the beginning of it.

They said: "All major browsers today employ coarse grain privilege separation to limit the impact of vulnerabilities. They run renderers, the portion of the browser that handles untrusted user content from HTML parsing, including JavaScript execution and image decoding and rendering, in separate sandboxed processes. This stops web attackers that manage to compromise the renderer from abusing local OS resources to, for example, install malware."

They said: "Unfortunately, this is no longer enough. Nearly everything we care about today is done through a website. By compromising the renderer, an attacker gets control of the current site, and often any other sites the browser has credentials for. With services like Dropbox and Google Drive, privilege separation is insufficient even to protect local files that sync with the cloud.

"Browser vendors spend a huge amount of engineering effort trying to find renderer vulnerabilities in their own code. Unfortunately, many remain, frequently in the dozens of third-party libraries used by the renderer to decode audio, images, fonts, and other content. For example, an out-of-bounds write in libvorbis was used to exploit Firefox at Pwn2Own 2018. Both Chrome and Firefox were vulnerable to an integer-overflow bug in the libvpx video decoding library. Both also rely on the Skia graphics library, which had four remote code execution bugs until recently.

"To appreciate the impact of these vulnerabilities and the difficulty of mitigating them, consider a typical web user, Alice, that uses Gmail to read email in her browser. Suppose an intruder, Trudy, sends Alice an email that contains a link to Trudy's malicious site, hosted on sites.google.com. If Alice clicks on the link, her browser will navigate her to Trudy's site, which can embed a .ogg audio track or .webm video to exploit vulnerabilities in libvorbis and libvpx and compromise the renderer of Alice's browser. Trudy now has total control of Alice's Gmail account. Trudy can read and send emails as Alice, for example, to respond to password reset requests from other sites Alice belongs to. In most cases, Trudy can also attack cross site, i.e., she can access any other site that Alice is logged into, for example, Alice's Amazon.com account.

"Recent versions of Chrome, and upcoming versions of Firefox, support Site Isolation, which isolates different sites from each other, for example, *.chrome.com from *.amazon.com, to prevent such cross-site attacks. Unfortunately, Trudy might still be able to access drive or pay or cloud.google.com" - because she has access now to gmail.google.com, so she could get drive.google.com, pay.google.com, or cloud.google.com - "which manage Alice's files, online payments, and cloud infrastructure, since the renderer that loads the malicious .ogg and .webm content might still be running in the same process as those origins." Thus process isolation doesn't help.

They said: "For many sites, Trudy might not even need to upload malicious content to the trusted victim origin, sites.google.com in our example. Most web applications load content, including images, fonts, and video, from different origins. Of the Alexa top 500 websites, for example, over 93% of the sites load at least one such cross-origin resource. And the libraries handling such content are not isolated from the embedding origin, even with Site Isolation." They conclude: "To mitigate these vulnerabilities, we need to harden the renderer itself. To this end, we extend the Firefox renderer to isolate third-party libraries in [what they're calling] fine grain sandboxes. Using this, we can prevent a compromised library from gaining control of the current origin or any other origin in the browser."

So as we've often observed on this podcast, our web browsers have become the largest attack surface that we routinely extend out onto the Internet. By their very nature, web browsing is insecure. It's inherently insecure. We're actively soliciting sites we know little about or even trusted sites that we may trust, but they themselves may have been compromised. Whatever. Whether an unknown site or a trusted site, they are sending tons of code to our browser, which our browser interprets, and asks our browser then to fetch massive third-party libraries and advertisements from all over the Internet, running code that we've never seen before in the context of our browser. We bring them all in, all of this stuff in to be processed, rendered, executed, and displayed. I mean, I'm so happy we have these guys, and guys like them, watching our backs. Lord knows we need it.

So this new technology is called RLBox, R-L-B-O-X. It will first be deployed in Firefox 74 for Linux, which is set to be released early next month, so a few weeks from now, in early March. Then the next month, in April, RLBox will ship for Firefox 75 for the Mac. And its development in Firefox for Windows will eventually catch up. So Firefox appears to be leading the industry as a result of this effort. And I'm sure once the technique has been proven and has matured, I'll be shocked if Chromium doesn't go sort of basically follow Firefox to this next step and employ these fine-grained sandboxes in order to protect our browser environment from the libraries that they run. Right now those libraries are running in process. They are not being separately sandboxed from the pages that they load. So that will be a good thing.

Also, just this morning, Chrome users running Chrome on Windows, Mac, and Linux were updated to 80.0.3987.122. It's not a big problem. I doubt it would affect anybody. I checked, and that's the Chrome that I had. It doesn't affect Chrome OS nor iOS or Android. So just the desktop platforms. It was released this morning to address three security bugs, including a zero-day vulnerability, which as its designation as zero-day implies was being actively exploited in the wild. We don't yet know anything about the attacks, only that it has a CVE tracker which describes it sort of generically as a type confusion in V8, V8 being of course Chrome's JavaScript interpreter compiler. We do know that the use of this bug in attacks was discovered exactly a week ago, on February 18, by a member of Google's Threat Analytics group that keeps watch on the things that Chrome encounters when its users are using it out on the 'Net.

This is, as I mentioned at the top of the show, the third zero-day that has been discovered in Chrome in a year. The first one was last March, so it's almost exactly a year ago that the first one was found. The second one was found in November, and the third one last week. So, you know, this is to be expected, since Chrome has become the Internet's number one browser. It's now the number one target.

So, SpinRite. It is where I am finally, happily, spending all of my time. I decided that in order to get back with the plan, I needed to go back and reread all of the newsgroup postings from 2013. I began in May of 2013. There were, well, if I eliminate the postings before and then the postings since the burst of work on 6.1, there were about 6,000 posts for the period that I was actively rolled up and working on 6.1. I've read about 1,200 of the 6,000, so I have about 4,800 remaining. And it has just been the most wonderful thing. The way I tend to operate in the newsgroups is to sort of journal what I'm doing, and what I expect to do next, and what I'm going to try, and then the results of that. I'm always producing incremental code for the gang in the newsgroup to experiment with and pound on, and they post their results. So I'm able to get a broad cross-section of hardware through time. It ends up being so useful.

But never before has there been a seven-year hiatus where I wanted to pick up where I was, but, you know, seven years. So I completely forgot what I was doing back then. And so this has just been incredibly useful for helping to bring me back up to speed. It just saves a lot of time because basically I have a detailed log of the whole beginning of this project that is still there, and I'm able to plow back through it.

One of the things that - oh, and I have been producing some new code, and the gang that's hanging out at the newsgroups has been running the code. I made a decision seven years ago which was the right decision then, and is not the right decision today, which was that the AHCI controller, which we all have to run our SATA drives, while the AHCI was a new spec, they had the option to operate in legacy mode. And, I mean, there are still motherboards where that's the case. And, for example, even Windows 10, I'm seeing posts out on the 'Net where people are asking, I installed Windows 10 when my motherboard or my system was set to legacy mode. I want to switch it to AHCI mode, but I tried it, and it won't boot. What do I do? Well, it turns out there are ways to get the drivers installed and then switch over, and Windows 10 knows about it because it uses different drivers.

Seven years ago it was much more expedient for me to know that legacy mode was present and to have SpinRite 6.1 use legacy mode. It is less sophisticated than AHCI. But AHCI is this huge, it's like sort of this next-generation controller for hard drives that definitely makes sense if you're on a server platform in a multitasking environment and, arguably, even on a personal workstation, where there's just a lot of stuff going on. But it didn't ever make sense for SpinRite, where basically you're in DOS, and SpinRite is saturating the use of the drive.

I mean, the things that AHCI provides, like very sophisticated command queuing, and you're able to chain jobs of things that you want the drive to do, in a chain of individual descriptor blocks. And the controller itself, it's basically a little microcontroller. It's able to read the descriptor block, see what it's being asked to do, and then go set the drive up and perform the transfer, reading this block of sectors into this block of memory. And once it's done, it marks the descriptor as completed and then chains to the next one and so on. So it's amazingly powerful. But it was just, like, overkill, totally, for SpinRite seven years ago. And it would mean that I was able to get 6.1 to its users - that was the original plan - a lot sooner, before SQRL, of course, came along.

Revisiting this today, it is no longer the case that legacy mode is ubiquitously available. It's obviously still in some places, but lots of the users who are using my PCI enumeration code, which I wrote first seven years ago and then updated, it's just no longer the case that it is available. And of course in order to fulfill my promise of having SpinRite run on contemporary systems, that means I have to support AHCI. So that, I realized a couple days ago, that decision from then had to be revisited and changed.

So essentially I was also surprised to see how much I got done during that four-month sprint from May through August of 2013. I have very mature code. I've got the flat real mode stuff all working and well tested. I'm able to allocate 32MB transfer buffers, not 32KB transfer buffers, and transfer 64,000 sectors at a time. So all of the legacy stuff will stay there because of course it will run on motherboards that are in legacy mode or in earlier older hardware that didn't even have AHCI. So all of that stuff will work with very large transfers, all bus mastering and operating the way we want.

And SpinRite will - the next thing I'm going to do is going to sit down and write AHCI support, add that to the platform that I already have. And then once we get that integrated into SpinRite, we're ready for the next release. I do have bad news on the Mac front. I know that older Macs did have a BIOS compatibility built in because I had SpinRite running on my MacBook, it was my MacBook Air, years ago. The reason SpinRite wouldn't run was that the Mac didn't emulate the keyboard hardware. It emulated the keyboard through the BIOS. The Mac uses a USB-based keyboard. And so it provided BIOS emulation, but not hardware emulation. When SpinRite starts up, it switches to reading the keyboard hardware, which is why the keyboard would freeze on people trying to run SpinRite on their Macs. That was easy to fix, and I will fix it so that, for Macs that do have that BIOS emulation, SpinRite will run on the Macs with no trouble.

The problem is Apple dropped the BIOS emulation in more recent Macs. They are UEFI only, and SpinRite is not going to run on them soon. SpinRite will eventually run on them because Intel has also announced their intention to drop BIOS support for future hardware platforms moving forward. So for now, because FreeDOS doesn't even begin to think about running on top of UEFI, no DOS runs on the EFI platform, that's a complete deal breaker. So it's clear that the only future SpinRite has is if it's also able to run on EFI. So I think I'm going to have to do that before I switch over to the plan for a complete SpinRite rewrite because EFI is coming, and it'd be nice to be there ahead of that, rather than always being behind the curve.

So anyway, that's where I am at this point. I've got still a lot more postings to read and catch up from the past. 6.1 will support AHCI, which was not my plan seven years ago, but times have changed since then. And I will have something that runs on current PC hardware. I don't know when. I never know when. But it's all I'm doing now, so I'm excited to be back at it and to be working on it. Just it makes me feel good, makes everybody feel good. So I'm glad for it.

Leo: Nice. Well done.

Steve: Something that did not make everyone feel good happened last week. Pause for a sip. The first I learned about it, before it was even in the press, was a letter that I received from Dean Taylor, who's the senior account manager at DigiCert. He said: "Dear Steve. Earlier today Apple announced that Safari will only trust certificates with a validity of 398 days or fewer, that is to say one year plus a renewal grace period." He said: "This policy goes into effect September 1st of 2020." And it's like, what?

So Apple announced last week, Apple announced that Safari will only trust certificates having a validity period, that is, the time from first valid, what do they call it, valid after and then not valid after, something like that. They have a very specific way of delineating when the certificate becomes valid and then when it expires. The point is that, if the span of those dates is greater than 398 for any certificates with a valid, a first valid date of September 1st or later, that is to say, issued after August 31st of 2020, if that span is greater than 398 days, it will not be trusted by any Safari on any Apple device - iOS, Mac, tvOS, iPad OS, anything.

So Dean goes on in his letter which is announcing this to me as a DigiCert customer: "Certificates issued before that date are not affected and do not need to be replaced or modified. You can continue to issue two-year certificates until August 1st, 2020" - and I think I will - "and use them until their expiration. This announcement was made by Apple," he says, "on February 19th at the CA/Browser Forum, an industry standards group meeting.

"While it's generally accepted that short-lived certificates will increase the security of the SSL ecosystem, we've been working with the browsers to time this change in a way that reduces the impact to our customers. While Apple's decision was unilateral, we already have tools in place to make short-lived certificate management easier, and we're working on additional solutions ahead of this change to offer you greater certificate lifecycle automation options."

He says: "I know this impacts your certificate management practices. That's why I wanted to let you know about the coming change and tell you that we are responding. At DigiCert we always put..." and blah blah. He goes into, you know, marketing speak. But he said, oh, here he said something: "Before the Apple changes occur, we'll add the ability for you to purchase multiyear certificate subscriptions to smooth planning and reduce the yearly work of buying and installing certificates. These subscriptions will let you reissue, renew, or replace a certificate as frequently as you need to without incurring additional fees. Our intention with offering subscriptions is to save you time and money."

And then he says: "Hand in hand with your multiyear certificate subscriptions is the ability to automate the entire certificate lifecycle. As the industry moves to shorter certificate lifetimes, automation is the key to keep your business running smoothly. To save time, avoid annual manual updates, and to avoid site downtime, DigiCert CertCentral offers you several ways to automate your SSL certificate needs: robust APIs, ACME integration, and our certificate automation tool. Automation allows you to spend more time doing what you want to do and less time managing certificates. If you have any questions," blah blah blah.

Okay. So remember, like, where we've come from. Back when this podcast was just beginning, when Honey Monkeys were crawling around, and come to think of it, when Leo's favorite password itself was "monkey"...

Leo: Monkey123, yes.

Steve: It was possible to obtain an SSL certificate from a certificate authority which would remain valid for eight to 10 years.

Leo: Oh, I wish I had.

Steve: Oh, from the time of its purchase.

Leo: Actually, it would have expired by now.

Steve: Yeah. Then, in 2011, the certificate authority, the CA/B, the Certificate Authority Browser Forum, CA/B Forum, which included all of the browser makers and the certificate authorities, decided that was too long, due to the, as we know, the enduring flakiness of certificate revocation and the possibility that a rogue certificate might escape and be honored for as many as 10 years.

Leo: Yeah. That's too long, for sure.

Steve: That's too long.

Leo: And just as a sidebar, we've talked about revocation for a while. There's no notion that it would be - it really isn't technically doable; right?

Steve: It technically is. It is.

Leo: You've always said it could be done; right?

Steve: Well, we have the technology to do it. It's called "stapling." The idea is that - so the problem with revocation is that OCSP, the Online Certificate Status Protocol, the OCSP servers have historically been flaky. And they could be DDoSed, for example. So the idea would be the browser gets a certificate. And it looks at it and goes, oh, I need to see if this is still valid, you know, if it's a valid certificate. It then makes a query to the OCSP URL that is in the certificate. The certificate says you are free to check the current status of this claimed to be and valid by calendar certificate at this URL. So the browser goes and queries the OCSP server. The OCSP server is run by the certificate authority that signed the certificate. And it says, yeah, that's valid right now. No revocation.

Leo: So the advantage of that is that's instantaneous revocation. The disadvantage is, and nobody I guess wanted to assume the cost of this check to OCSP every time you check the cert.

Steve: Yes. The temporal cost of...

Leo: Yeah.

Steve: I mean, because even that, even that slows down the presentation of the page.

Leo: And everybody wants pages to load instantaneously.

Steve: Yes.

Leo: So this would be a preferable situation, but nobody's willing to do it. No browser is going to slow it down that much.

Steve: Well, but there's more, Leo. Then, to solve that problem, the really cool innovation known as "stapling" occurred. With stapling, the web server which is sending out the certificates to be trusted, it periodically gets a signed OCSP statement and staples it to the certificate.

Leo: Oh. So it does that behind the scenes.

Steve: Yes.

Leo: Doesn't slow you down.

Steve: Exactly. So the browser receives the certificate and a recently updated affidavit signed by the certificate authority saying, yes, we are reasserting an hour ago that this certificate is still valid.

Leo: What would the typical stapling window be?

Steve: It could be a few hours. The idea is that, as the server begins to see that its current certificate attestation is getting near the end of its life, like a few hours, it starts asking for an update from the certificate authority's OCSP. It updates it, staples it, and then sends the certificate.

Leo: So why didn't we do this?

Steve: I know. Once upon a time it wasn't widely available. All the web servers now support OCSP stapling. So - I know. The problem has been solved.

Leo: So this would have been preferable because, instead of having a three-month or a one-year expiration, you could revoke a certificate within a matter of hours.

Steve: Correct.

Leo: So far preferable, from a security point of view.

Steve: Even better than a one-year expiration.

Leo: Because now a phony certificate could be as good as a year long; right?

Steve: It could. It could. Right. Right. And that's why also Apple just sort of, I don't know, that's why I think there was probably more going on behind the scenes.

Leo: Yeah, politics.

Steve: Yes. I looked for - this is, like, every - this is a big deal. Basically this cuts maximum certificate life in half. It had been two years. Now it's one year. And it was not done by consensus. It was done unilaterally by Apple. They've got 17% of the browser share. No webmaster is going to have a certificate that no iOS or Mac user can trust, you know, that brings up warnings that the browser won't show them. So, I mean, Apple was big enough to pull off a unilateral coup, essentially.

So as a consequence there's a lot of fur flying. I looked for something definitive, and I found it, again from my favorite certificate supplier. This is from Dean Coclin, C-O-C-L-I-N, the Director of Business Development. And he did a posting that was titled "DigiCert's Position on One-Year Certificates." And it provides some additional background information that I thought our listeners would find interesting. And it's not hype.

He said: "At the CA/Browser Forum in Bratislava, Slovakia this week" - and that was last week - "Apple announced that beginning September 1st, newly issued publicly trusted TLS certificates are valid for no longer than 398 days. This followed a long history of the CA/B Forum community working to reduce certificate lifetimes and improve security, while balancing the needs of business owners in transitioning to shorter validity certificates.

"In August 2019" - okay, so that was last summer - "CA/B Forum Ballot SC22 was introduced by Google to reduce TLS certificate validity periods to one year. CAs reviewed this proposal with their customers and produced thousands of comments from users, which mostly showed opposition due to the additional work required by IT teams to handle shorter validity periods. The ballot failed in the Forum, which meant certificate maximum lifetimes remained at two years.

"At one time, certificates were offered with a maximum validity of three years. A few years ago they were reduced to two. Fast-forward to this week's Apple announcement, which ultimately does what Ballot SC22 failed to do: reduce certificate lifetime to one year." He asks rhetorically: "Why did Apple unilaterally decide to enforce a shorter certificate lifetime? Their spokesperson said it was to 'protect users,' he has in quotes. We know from prior CA/B Forum discussions that longer certificate lifetimes proved to be challenging in replacing certificates in the case of a major security incident. Apple clearly wants to avoid an ecosystem that cannot quickly respond to major certificate-related threats." But again, Leo, your point is perfect, and that is, okay, so it's still a year.

Leo: That's a long time.

Steve: You could do a lot of damage in a year.

Leo: It would have been a lot better to put the energy behind OCSP and stapling.

Steve: Yes. Yes.

Leo: In my opinion, your opinion.

Steve: Yes.

Leo: It's bizarre, to be honest. Google does this, too. I mean, they did this with TLS. They do these kind of...

Steve: Yes, unilateral.

Leo: Unilateral.

Steve: We're going to throw our weight around.

Leo: Right. We're powerful.

Steve: Because we - yeah. So who knows, like, what politics was going on behind the scenes. My guess is, as I said at the top of the show, I bet not everybody was surprised by Apple's statement. I'll bet you the other browsers, they were kind of hobnobbing and like, you know...

Leo: They may have drawn straws. Apple might have actually gotten the short straw. Like it could be Firefox, Google, Apple all got together, look, somebody's got to do this. I've got three straws. Short straw announces. Because any one of them's big enough market share that it forces the hand. So maybe Apple didn't want to do this. I don't know. We just don't know.

Steve: Yeah. He said: "Apple clearly wants to avoid an ecosystem that cannot quickly respond to major security-related threats. Short-lived certificates improve security because they reduce the window of exposure if a TLS certificate is compromised. They also help remediate normal operational churn within organizations by ensuring yearly updates to identity such as company names, addresses and active domains."

Leo: That's fine.

Steve: That's kind of a good point.

Leo: And Let's Encrypt is three months. And if you're using a Let's Encrypt cert, you don't have to worry about this. Problem is we use wildcard certs. And so I can't use Let's Encrypt. I have to go through a fairly elaborate rigmarole manually. I guess we'll have to automate it. Do you think they'll go shorter than a year? Like this is the first in a series?

Steve: I think this is.

Leo: Yeah.

Steve: I think, yeah, I think that this all - it doesn't really force automation. But, like, I mean, there are, I mean, as I told our listeners a couple weeks ago, I recently decided, okay, I've got all kinds of subdomains, you know, sqrl.grc.com, blog.grc.com. I'll have a spinrite.grc.com. I mean, I'm liking that. I was also liking having EV certs, but they don't allow wildcards.

Leo: And those are gone anyway now; right? I mean, they've gone away.

Steve: Well, and they no longer show you any little extra bling in the browser's URL. So it's like, okay.

Leo: So what's the advantage?

Steve: Yes.

Leo: So somebody's saying Let's Encrypt does do wildcards. See, maybe we should just go to Let's Encrypt because that's free.

Steve: Yeah.

Leo: Maybe we should just go to Let's Encrypt. We just unfortunately bought two years for both twit.tv, you know, *.twit.tv and *.techguylabs.com. But in a couple of years maybe we'll just all go to Let's Encrypt.

Steve: Well, so I think we're going to see automation. We heard them say, you know, because first I was thinking, wow, you know, not only does it double the work on the user side, it doubles the work on the issuer side, if they're going to have to reverify this twice as often. And then I realized, oh, that's this idea of having a subscription. The idea is you could, I mean, that gives them some lock-in, which they like. You're subscribing for a longer period of time. I'm sure you get a better rate per year for the more years you subscribe. Then on some sort of schedule they will reverify your identify. But as long as you have a subscription, then you're able to reissue your own certificates.

And frankly, that's the way DigiCert already works. When I reissued my own hybrid non-EV *.grc.com, which is what I did, and www.grc.com and grc.com, when I reissued that, I did it myself in the middle of the night because DigiCert already has the automation where they know how old their certification of me is, and they periodically call up and say, hey, just want to make sure you're still answering the phone and you haven't moved and blah blah. I say, oh, yeah, still me. Sometimes they'll send me a file that I have to stick on the root of GRC to re-prove my ownership of the domain.

Leo: Yeah, we have to do that kind of thing, yeah. In fact, we're about to do that. We've got the CSR and the whole thing.

Steve: Yeah.

Leo: So why not use Let's Encrypt? Shouldn't we all, I mean, is there a - it's free; right? I pay - you and I pay a lot of money for these things.

Steve: Yes. I still sort of think that - okay. So the problem is there is massive fraudulent use of Let's Encrypt. Every possible misspelling of PayPal that you can imagine...

Leo: Why wouldn't they?

Steve: ...has a certificate now from Let's Encrypt. And I think that browsers, you know, they used to show us what was secure and what wasn't. Then they went away from that, where we're assuming security. They used to show us what was EV and what wasn't, and then they went away from that. I'm betting someday they're going to show us whether the certificate was issued without verification.

Leo: Oh, automated or not, yeah.

Steve: Yes. Because what you and I are getting from wherever you're getting yours, I'm getting mine from DigiCert.

Leo: We were DigiCert, and then Russell said, you know, it's about half as expensive at GoDaddy. And as much as I hate GoDaddy, it's still 800 bucks for two years. So, you know, it would have been more.

Steve: Yeah. So what Let's Encrypt, the only thing they can issue are DV certs, Domain Validation. I'm using and you're using OV certs, Organization Validation.

Leo: Oh, yeah, that's right.

Steve: And so I'm just kind of thinking, at some point, as automation happens, as anybody could, I mean, anybody anywhere can get a domain that is a spoofed domain that's PayPal misspelled somehow, and now have a security certificate for it. No serious certificate authority worth their salt would issue a certificate for a domain clearly intended to spoof a valid, well-known domain like PayPal. So that's sort of the last vestige of integrity is was there a human in the loop at the certificate authority who said, "Yes, you guys are real." And at the moment, that isn't shown. But I have a sense in the future it'll be a means of giving the user some additional sense because once the world goes HTTPS, even the malware all is now HTTPS because they're all using Let's Encrypt.

Leo: Yeah.

Steve: And DigiCert is using ACME. ACME is that automation protocol.

Leo: That's what LE uses, as well; right? They use that, yeah.

Steve: Yes, yes, yeah. Although they're backing it up with - they're able to issue organization verification certs.

Leo: Right. That's what I want. Right.

Steve: Because, yeah, I really think you do.

Leo: Yeah. It's only 800 bucks a two-year cert, so it's not awful.

Steve: Yeah. And so I don't think that - they're certainly not going to charge $800 for a one-year cert. They'll probably...

Leo: Yeah, sell me two one-year certs.

Steve: Well, or say, hey, do you want to, are you willing to commit to stay with us for 10 years? If so, we'll give you a discount, I mean, just like domain names are. We'll give you a discount if you buy a longer period of time. And during that time, you'll be able to reissue, and in fact you'll have to reissue, at least every year. But you'll be able to reissue any time you want to within the validity period of that subscription. So basically we're changing the way we interact with certificate authorities. And once that's happened, Leo, as you forecast, I think we'll be dropping to even more often than once a year because we'll have the mechanisms in place. But again, we solved the problem with stapling. It's a perfect solution. And it's like, oh, well, yeah, hmm. We're not going to do that.

Leo: If I were DigiCert, I'd be pissed because it does feel like maybe they're going to be putting them out of business. Like the goal is, well, maybe not. But, I mean, honestly, a lot more business will go to Let's Encrypt.

Steve: They're putting pressure. Yes, yes. Although also it turns out these certificate authorities are generating certificates for all kinds of other things. Enterprises have really expansive needs for authenticating endpoints, and things that, well, remember you can't get a certificate from a CA for a non-public domain. So enterprises have all kinds of other needs for certificates for authenticating their own internal VPNs, all their internal connections. And apps are often signed with certificates that are only recognized by enterprise. So there's all kinds of other uses, too. But anyway, Apple just said...

Leo: What a surprise. Apple just did this, yeah.

Steve: Yeah. We're us. We're just going to force it.

Leo: Yeah.

Steve: Even though just six months ago a consensus ballot failed.

Leo: I've got this new conspiracy theory. They got together, and they drew straws.

Steve: Spin the bottle, yeah.

Leo: Apple got the short straw. You're going to announce it. We all want to do it. You're going to announce it. And then it'll be forced. The issue is forced. Okay. Thank you, Steve.

Steve: I'm going to get back to work on SpinRite.

Leo: Yay. Details at GRC.com. That's where you'll find, not only SpinRite, the latest version. If you buy a version today, by the way, you'll get the new version when it comes out. So don't worry about upgrading.

Steve: Actually, if you buy it today, you get all the pre-release versions before it comes out.

Leo: Oh.

Steve: Because you'll be able to use your serial number or transaction code in order to get access, pre-release access.

Leo: Nice.

Steve: So there's a little spiff there, too.

Leo: Yeah. I bet there's a few people using it who have not yet paid for it. So here's an opportunity to make yourself right with the world and get some nice benefits. GRC.com. Steve has lots of free stuff, too, including ShieldsUP!, his famous, world-famous routing checker. If you listen to the show, you might want to get it there. He has 16Kb versions, that's the smallest file size, and English-language human-written transcriptions, which is great if you like to read along while you listen. That's all at GRC.com.

We have the show, as well, 64Kb audio, hundred whatever video, big. And you can get that at TWiT.tv/sn. It's also on YouTube. Best thing to do is subscribe in your favorite podcast application. You can get it that way. Just search for Security Now!, or better yet, search for TWiT. Subscribe to all the shows you want, and then they'll automatically be downloaded, and you'll never be without stuff to listen to on your device. We record the show every Tuesday, 1:30 Pacific, 4:30 Eastern, 21:30 UTC. If you want to stop by and watch you can, or listen, audio and video streaming at TWiT.tv/live. If you're doing that, join the chatroom at irc.twit.tv.

We had a little outage yesterday. The chatroom went out, the wiki went out, and most importantly from our point of view, the RSS feeds went out. Our server company, Contegix, had a switch failure. They rebooted and failed again. It took them about a couple hours, five hours something, to get back. So if you were a little slow getting the downloads yesterday, or maybe couldn't get in the chat, that's why. But it's all okay now.

If you're going to be at RSA, we will see you tomorrow. Don't forget to email tickets@twit.tv or go to TWiT.tv/blog to get the secret password for the fabulous LastPass party. Steve is on the Twitter at @SGgrc. That's a good place to leave questions for him. He takes DMs from anybody, and he tweets there regularly, too.

Steve: And speaking of LastPass, I have confirmed my participation in the next event with you, Leo.

Leo: This fall. We're going to do an event in San Francisco. And everybody will be invited to that. We'll give you details to follow.

Steve: Yeah.

Leo: Thank you, Steve. I can't wait. You and Lorrie coming up? Dinner's on me.

Steve: It'll be great, yup.

Leo: We're going to have a lot of fun.

Steve: We will indeed.

Leo: Yeah, there's good restaurants in San Francisco, I hear. I hear that.

Steve: Is it actually going to be in the city? All I knew was the Bay Area.

Leo: I think so. I don't know. You know what, ask Laura. She's in charge. I should never talk out of school because I'm the last person who knows. Thank you, Steve Gibson. Have a great week, and we'll see you next time on Security Now!.

Steve: Right-o.