Yahoo & Primal Worries

Description: Leo and I discuss today's Windows Update changes for 7 and 8.1. An exploit purchaser offers a $1.5 million bounty for iOS hacks. WhisperSystems encounters its first bug. An IEEE study reveals pervasive "security fatigue" among users. We've got Firefox and Chrome news, WoSign Woes, Samsung Note 7 news, some errata, a bunch of miscellany, and a look into new Yahoo troubles and concerns over the possibility of hidden trapdoors in widely deployed prime numbers.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-581.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-581-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve and I renew our acquaintance after a few weeks gone. There is a lot of news, lots of things to talk about, including our own personal favorite shows and that kind of thing. We'll also get into the Yahoo hack and what Yahoo has done, which is frankly reprehensible. And Steve will explain a new problem, maybe you read about it, the trapdoor prime issue, and why it may not be a bad as it sounds, although something to keep in mind. Anyway, details to come, next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 581, recorded Tuesday, October 11th, 2016: Yahoo & Primal Worries.

It's time for Security Now!, the show where we cover your security and your privacy online. And the king of security and privacy is here, Steven "Tiberius" Gibson of - this is now your middle name, by the way - of GRC, Gibson Research Corporation. Hi, Steve. I missed you. Thanks to Father Robert for filling in.

Steve Gibson: Likewise. Well, we did, you know, Father Robert is a great co-host, but this feels like home, having you back.

Leo: Well, we have been, you and I, together for 11 years doing this show, and many years before that. So it is kind of just a couple of buddies, sitting around over a cup of coffee, cup of Joe, talking security.

Steve: Comfortable old shoe.

Leo: That's me. That's me. A little more comfortable after living on a cruise ship for two weeks, I might add.

Steve: I would imagine so. And nobody got sick. Everybody's healthy.

Leo: Everybody's great. We had a great time. It was a really fun trip. And I have to say, you know, towards the end I started doing shows, having imaginary conversations with people and stuff. And especially when you read security stories or something, I go, oh, I'd love to talk to Steve about this one. Or we're in Russia, and we're thinking about Putin.

Steve: Putin and hacking email and so forth.

Leo: Hacking and, oh, I'd love to talk to Steve. I'd love to know what he thinks. So now I get to find out what Steve thinks about a lot of stuff. Weren't we - we're going to cover Yahoo; right?

Steve: Yes. In fact, the title of today's show is "Yahoo & Primal Worries."

Leo: Oh, whoa, well, sounds deep.

Steve: Because a paper was just released which brings into question some assumptions which, 25 years ago, when some of the current technology was put in place, were questioned. And then there was a little controversy, but it was dismissed. However, some researchers in France and the university of something in the states, I can't remember what, they've successfully demonstrated an implementation or an exploit of the fear that was originally voiced 25 years ago. And it's not good.

Of course we have to talk about Yahoo because just, like, could - in fact, there's even some news some people may not have caught up with yet because it's just happened, which is like yet another nail in the coffin. And it's funny because last week Father Robert and I were discussing this. And in one of the Q&As that we did, sort of the person posed the question, you know, given what we know, should I move away? And it was like - and we came down on saying, how could you not? I mean, if you actually care about security. I mean, for convenience, if you're there, okay, fine. But if you truly are giving more than just lip service to security, yeah, you have to move. Anyway, and so there's more reason to feel that way even than there was last week.

But a lot of stuff is going on. We're going to talk about this being the October Patch Tuesday, the first one. This is the one where everything changes. And I fired up my Win7 machine an hour early to give it a chance to get itself synchronized; and, sure enough, no more batch of patches, just a monolith. So we're going to talk about that.

There's a $1.5 million bounty for iOS hacks. And in the coverage of this, there was some interesting commentary about why Android hacks receive less money than iOS hacks. WhisperSystems had its first bug discovered in attachments for its Signal protocol. The way they handled it we'll talk about. An IEEE study that wasn't intended to yield the

results it did surprised them, and that's being presented, and we'll talk about the results. A little bit of Firefox and Chrome news. The ongoing woes of WoSign drama. I wanted just to briefly mention the Samsung Note 7, which has been much in the news also. And then we've got a little bit of errata, a bunch of miscellany.

And then we're going to wind up talking about catching everybody up on the latest Yahoo troubles, and take a look at what this concern is about prime numbers. As usual, the headlines and the first paragraphs of the press coverage tend to be probably overwrought, which is not to say there isn't a concern. But this podcast is all about helping our listeners to sort of calibrate these stories and what they actually mean in the real world. So I think, for 581, another great podcast.

Leo: Another jam-packed podcast. I'm disappointed to hear that everything's not settled down since I've been gone.

Steve: Well, you know, Father Robert and I were able to do a bunch of Q&As.

Leo: I'm glad about that, yeah.

Steve: Because you and I hadn't been because there was just too much to talk about. Like every single week was a cornucopia.

Leo: This Yahoo thing, I can't, I just can't wait to hear what you have to say about it. All right, Steve. On with the latest.

Steve: So our Picture of the Week on the first page of the show notes is just sort of keeping an eye on what Let's Encrypt is doing. There was some coverage that was claiming that Let's Encrypt had now become the largest certificate authority on the Internet. And it's sort of a function of how that's measured because there's numbers of certificates issued. There are the numbers that are currently in circulation and so forth. But anyway, this is a graph which shows, maybe for the first half of this year, pretty much an exponential curve.

Leo: Wow, wow.

Steve: Now, from the summer until now, with the exception of a little blip, it's gone sort of linear. But it's linear at a good clip, like a million certificates every couple months. So it's clear that...

Leo: So those numbers on the left, that scale's not labeled, but that's millions.

Steve: Those are millions, yes.

Leo: Wow, wow.

Steve: Although I've seen the number six million, so I think this chart must be those currently in use because...

Leo: Each month, or maybe it's not cumulative, yeah.

Steve: Well, the Let's Encrypt certs, also there may have been people experimenting with them, or revoking them, or having short expiration times. So there's a lot more churn. Which it makes sense to do because, remember, one of the ideas that we've discussed in the past for improving just the general robustness of the CA system is to shorten the lifetime of certificates. Right now they're two to three years, typically, depending upon what type of cert. But what that means is that, if a private key were to escape, the proper solution is revocation. But we know that the certificate revocation system is badly broken.

So the alternative is to issue very short life certificates so that even a stolen one would self-expire, that is, the signature on it would take it out of service quickly anyway. The problem is that's not practical when there's a lot of per-certificate issuance overhead. So the fact that the Let's Encrypt system automates and completely makes that whole process transparent, it means that it's completely feasible to roll or rotate your certificates very frequently, obsoleting them and having the automated system replace them with refreshed and timestamped signatures, but with none of them having a long lifetime.

So that's why it's important to draw the distinction between the total number of certs issued, especially in a model which encourages short lifetime certificates. Make that separate from how many are actually in use at any given time. But no matter how you look at it, it's been a smash success. And in fact, as I will show toward the end of the show, it is arguably this success, and in general the "going dark" problem, which drove the U.S. government to do what we believe it did with Yahoo. So we'll discuss that.

But it is Patch Tuesday. And as I mentioned at the top of the show, Microsoft did release a single monolithic bundle. And so essentially what we're going to get now is what we've traditionally called "rollups," where a bunch of individual patches were rolled up into one. It is more efficient because many times different patches are containing different versions of the same component. In Windows, sort of the unit of a component is an EXE, a SYS which is typically a device driver, or a DLL. And many times, three or four patches in a given month will each be making their own change to the same, to one of the same DLLs. So you're getting that same DLL four times.

And then the nightmare that I really salute Microsoft for even attempting is how do you allow users to selectively decline individual updates and have, like, everything else still work? It's amazing to me that they were able to do it for as long as they have, given the complexity of Windows, which is beyond any individual's comprehension at this point.

So what they're saying is they're going to do, essentially, with 7 and 8.1, starting today, what they've been doing with Windows 10, which is, on one hand, you can say taking away user choice. The flipside is, if there's just one set of updates, it can be more robust, better tested, and probably overall more reliable at the cost of no longer allowing users to pick and choose among them.

This, of course, as we've talked about, was very useful during the whole Get Windows 10 period because a number of the techier users were going in and, every single month, turning off the GWX reattempt to install itself and saying, no, I don't want that. No, I don't want that. That kind of thing won't be possible in the future. You simply won't be able to pick and choose. So today's update is important. It sounds weird even to say "today's update" because...

Leo: Hasn't been one update in years, yeah.

Steve: Oh, it's been 18. You know, we go, oh, my god, it's 18 updates that are fixing 57 different vulnerabilities. Now it's, okay, it's an update. There were five zero-day flaws, however, which have been fixed in this. They're each in IE and Edge, so each of the browsers has one. There's a zero-day in Microsoft Office. There's one in the GDI+ module which is exploitable through the web browser. And then the Internet messaging component of Windows has one. So, and of course, as we know, zero-day is not just nobody knew about it because typically these are all a surprise. It's that they learned about it by seeing it in use. So these are actively being exploited. So you're going to want to install today's update.

Now, in digging back into this a little bit, I thought it was interesting that what Microsoft has announced is that there will be essentially three monoliths from now on. There's what they call a "security only" update, which will be a single update containing all new security fixes for that month. Okay, so that's sort of an incremental - that's still not going back any further, but it's one thing. So all that's essentially doing is it's removing the granularity. It's a single update containing all new security fixes for that month.

Then there will be what they're calling in the notes, they said, that is, in their post, a "Security Monthly Quality Rollup." But I'm looking at my machine, and it calls it "October 2016 Security and Quality Rollup." Oh, now, here's one for the .NET framework, and then they have "October 2016 Security Monthly Quality Rollup" for Windows 7. So those are the two. And then they will be also offering a preview of the monthly quality rollup on the third Tuesday of the month, which is to say next Tuesday they'll be making a preview of what they're planning to offer two weeks later, on the second Tuesday of the next month.

So that's what's going on. Essentially, users who've never been intimately with Windows Update won't notice any difference. It'll just be like, okay, you know, update me. Oh, only one? Okay, fine. Those of us who know better realize that they've all been crammed together into a single blob. I guess I'm of two feelings about it. It will be very nice in the future for people who want to set up new older versions of the OS because right now these rollups only deal with current fixes. That is to say, and I mentioned it last week, everybody should make sure their Windows is current by today because then the single blob moves it forward.

But in coming months, Microsoft has stated they intend to reach back further and further in time, eventually all the way back to, for example, in the case of Windows 7, to Service Pack 1, which was the first and only official service pack, and you can get from Microsoft an image of Windows 7 SP1. And then the idea would be, once the monthly "quality rollup," as they call it, is comprehensive, a new Windows 7 install would be installing the old Windows 7 SP1 image, and then one single blob which would be the most recently issued monthly quality rollup, and you'd be done.

So I think this is good in the long term. We'll sort of have to see how it plays out. Again, if there were any problems caused by the granularity that they had traditionally been offering, while we lose that, if we get more stability, I think that's probably a net win for everybody, for Microsoft and for users.

I guess it was last week I saw that the exploit purchasing company Zerodium had tripled the price that it was offering, or the bounty that it was offering for iOS exploits, and doubled the price that it had been offering for Android. And Dan Goodin of Ars Technica, he had a nice little bit of coverage of this. He said: "A controversial broker of security exploits is offering $1.5 million for attacks that work against fully patched iPhones and iPads, a bounty that's triple the size of its previous offering.

"Zerodium also doubled, to $200,000, the amount it will pay for attacks that exploit previously unknown vulnerabilities in Google's competing Android operating system; and the group raised the amount for so-called zero-day exploits in Adobe's Flash media player to $80,000 from $50,000. After buying the working exploits, the company then sells them to government entities, which then use them to spy on suspected criminals, terrorists, enemies, and other targets."

Dan writes: "Last year, Zerodium offered $1 million for iOS exploits, up to a total of 3 million." And it paid it out. It paid three of those a million dollars each, after which it dropped the price to half a million dollars. On Thursday of last week Zerodium's founder - and I cannot pronounce his name - Chaouki Bekrar said, "The higher prices are a response to improvements the software makers, Apple and Google in particular, have devised that make their wares considerably harder to compromise." So that's the good news.

And then Dan quoted him a little further, explaining the price difference. He said: "Prices are directly linked to the difficulty of making a full chain of exploits, and we know that iOS 10 and Android 7 are both much harder to exploit than their previous versions." Asked why a string of iOS exploits commanded 7.5 times the price of a comparable one for Android, he said, "That means that iOS 10 chain exploits" - meaning front to back, soup to nuts, drop this on the phone and it just takes over the phone, a chain exploit - "are either 7.5 times harder than Android, or the demand for iOS exploits is 7.5 times higher." And he said, "The reality is a mix of both."

So as we know, Apple's had a focus on security. They've implemented much of their solution in the hardware architecture of their device, which is where I think they are able to arguably claim they've managed to really make it more difficult because they've locked so much up into proprietary hardware. And also that's the phone that high-value targets are using. And so governments and other agencies willing to pay an incredible amount of money for these things have the money and do. Wow.

Moxie at WhisperSystems reported the first bug that had been found of this type in the Android implementation - and it's important to say only the Android implementation - of Signal. And it was a wonderful bug. That's really the thing that brought it to my attention.

Leo: A wonderful bug.

Steve: It's so perfect. But get this. When the Android code retrieves an audio, video, or image attachment, it verifies a cryptographic MAC, as we know, a Message Authentication Code, to ensure that the attachment has not been modified in any way

while in transit. Two guys, security researchers, pointed out that a 32-bit integer was used to represent the attachment's length in that calculation.

Well, okay. We know what 32 bits is. That's the size of the IP address space. That's 4.3 billion. But that's 4GB. Which maybe once upon a time was unlikely. But, you know, it's a large attachment. But what this means is that, if the attachment size is greater than 4GB, the integer representing the attachment's length wraps back around, starting from a value of zero. That is, what it really needs is 33 bits, or 34 bits. Because once you go to all ones, and that is the maximum non-signed scalar value that 32 bits can hold, you add one to that, and the thing goes back to zero. It wraps around. That's why we call it a "wrap."

So if an attacker were to hack a Signal server - that is to say, they'd need to establish a man-in-the-middle presence somehow, and so hacking a server is one way to do it - and were to append 4GB of data to a legitimate, for example, 1MB attachment, while it was in transit from one end to the other, the code that verified the integrity of the attachment on the recipient's end would only see that initial pre-appended 1MB of data, the 32-bit value would show a 1MB length, even though it had tried to show a 4GB plus 1MB length. But because of that 32-bit wrap, the extra 4GB would be hidden, essentially, from it.

So, now, what this means is that it turns out it's difficult to exploit. The attacker would be forced to always append exactly 4GB, that is, they don't have any choice of the size of what they append. And Android will separately reject most media attachments that are that large. But they immediately recognized that this was an oops. They fixed it, pushed out the fix for Android and noted that exploitation is difficult; but it was wrong, and they have fixed it, and none of the other platforms were affected. So probably this was just a tiny mistake in the coding or the casting of the size of a value or something. Who knows what really went on behind the scenes.

I was curious that neither the iOS implementation nor their desktop flavor were affected, only the Android versions. Oh, and nor were any other consumers of the Signal protocol, for example, like WhatsApp over in Facebook and so forth. So the blog over on WhisperSystems said this is the first time that anyone has ever found a bug like this in Signal. So huge thanks to the researchers for helping to further improve the security and stability of the app. So not a big deal, but WhisperSystems handled it as well as they could.

Leo: Yeah. There's no such thing as a bug-free program.

Steve: There was an interesting article, and it didn't make it into my notes, but it was the idea of provably correct software. There are people working on it. And remember, bugs, you know, we spend half of our time on this podcast talking about mistakes.

Leo: All security holes are bugs, essentially; right?

Steve: Yeah, but it's also math. And we have seen subtle defects in Intel processors that made lots of news. There was a famous division error in one of the early Pentium chips that caused spreadsheet flaws, and then you were able to...

Leo: The 386, yeah.

Steve: Exactly.

Leo: It was like a rounding thing, yeah.

Steve: Right, right. So it is possible that the hardware has a problem. Again, a bug. But ultimately it's math. So, and this is what used - I used to have a bigger problem with this than I do today because 11 years of this podcast has educated me as much as it has all of our listeners about, I don't want to say the futility, but the challenge, certainly. And where we've done down is that the best we can do is responding as quickly as we can and try not to make mistakes. But if you do, fix it quickly and openly and move on.

Leo: Some day it'll be fun to do a show about mathematically provable, provably correct. Because I can't remember who I had the conversation with. I think it was on Triangulation. But they are endeavoring to create languages that are provably correct. Certainly if you're NASA, and you're launching a rocket that you can't modify once it gets out there, you want to get it as good as you can.

Steve: Well, and a perfect example was the cost of the space shuttle's code because it was very little code, and it was - pardon the use of the term - "astronomical."

Leo: In its potential for screwing up.

Steve: Oh, I mean, they absolutely had to get that right. And so the point is they did. But with the tools they had at the time, it was pure manpower, and how many eyeballs could look at it. And so the goal of the whole open source movement was, oh, we're going to have everybody looking at this code. Well, history demonstrates that four people look at it, three of them who didn't write it, who assume it's correct. And they kind of, you know, they're tired, and their coffee ran out a few hours ago and, oh, yeah, this looks fine to me.

Leo: Looks good to me.

Steve: And it just goes right past them.

Leo: But there is a mathematical - there is a field of mathematically provably correct software.

Steve: Yes.

Leo: Which boggles my mind that that even could be such a thing.

Steve: Well, we know that every time you add two plus three correctly, you get five.

Leo: Right.

Steve: No matter, I mean, we got that. The problem is we haven't figured out how to scale to something way bigger than that and have the same...

Leo: Given that a program is provable, you're right.

Steve: Yeah. So, and one of the problems is, you know, notice that what's happening with this crazy amount of storage and computation power is suddenly, kind of out of nowhere, AI is in the air again.

Leo: Yeah, yeah.

Steve: I mean, it's like, whoa, how did this happen? How does this thing actually know what I'm saying? So my point is that we have had an escalation in power over the last few years. Well, one of the other ways that could be applied is for this kind of software integrity. That is, right now we're writing software the way we have for 50 years, where every byte counted - well, okay. We're not writing it that way anymore. But still, the idea is it's an individual who's expressing in a non-natural language, algorithmically, what they're trying to get the computer to do. But there's a language gap, essentially.

Well, one of the other ways we could apply this crazy amount of excess power that we now have is to fixing that language gap, to fixing this problem of what the programmer intends not being expressed correctly to the computer. And a perfect example is, okay, why is it that on Android that attachment size was 32 bits? How did that happen? It would be really interesting to look closely at that. And we have, because this kind of thing fascinates me, in the past we've looked at exactly how it was that this mistake occurred, where somebody went in, and they had too many nested levels of parentheses, and they just got it wrong. And only when you really look at it you go, oh, look, these parentheses encompass this. And so the cast on that expression is done this way here and that way there, and that difference can be exploited. I mean, we're in the weeds.

And there hasn't been any dramatic change. There have been some efforts to create more abstraction in the way we express what we want the computer to do. I would argue that that'll be one of the next things to happen. At some point in the future, and it's probably a ways away because we know how slowly these things change, the way we program may be very different. Because it is just math. We should not be having these problems. I think it's clear that it's just inertia. On the other hand, I like to program the way I do.

Leo: You know, it's that debate of art versus science. Art is imperfect, and it's a lot

more fun for humans. And I think programming is, in some degree, an art as well as a science. But that's where the mistakes creep in.

Steve: You could argue that, for example, music could be programmed so that every note is precisely hit at exactly the right...

Leo: Yeah, and you wouldn't enjoy it.

Steve: Right. It loses its warmth and its something ineffable that is there that is valuable. And I don't know whether tubes actually sound warmer than transistors. I'd have to get into that one [crosstalk].

Leo: That's something else. That's another [crosstalk].

Steve: Okay. So the IEEE - I'm sorry. The NIST, the National Institute of Standards & Technology, the NIST conducted a survey where they were not looking for what came out. And it wasn't a huge survey. I saw an N=40. So they interviewed - they gave 40 typical computer users a questionnaire. But a surprising outcome is being reported and submitted to the IEEE's IT Professional Group. And that is they identified something known as - and everyone's going to understand this - security fatigue.

And in their own reporting of it, they said, "Security fatigue came oozing out of the questionnaires. People are tired and fed up with the burden of being made so responsible for their online security." And when you think about it, that's what we talk about all the time. When Jenny finally started using LastPass, she's like, oh, my god, how have I - and she was mad that I hadn't told her about it sooner. You know, "Why didn't you tell me about this sooner?" And I have another friend who still isn't using it, and he's got crazy passwords. And I say, "Mark, you're not using LastPass yet?" "No, no, not yet."

But the point is that, unfortunately, the way the system has evolved is the "solutions," unquote, that have been put into practice so far mostly make the user increasingly responsible for their security. That's the way the system works at this point. So in the abstract of their report they wrote: "Security fatigue has been used to describe experiences with online security. This study identifies the affective manifestations resulting from decision fatigue and the role it plays in users' security decisions. A semi-structured interview protocol was used to collect data," and they say N=40.

"Interview questions addressed online activities; computer security perceptions; and the knowledge and use of security icons, tools, and terminology. Qualitative data techniques were used to code and analyze the data identifying security fatigue and contributing factors, symptoms, and outcomes of fatigue. Although fatigue was not directly part of the interview protocol, more than half of the participants alluded to fatigue in their interviews. Participants expressed a sense of resignation, loss of control, fatalism" - just shoot yourself.

Leo: Throwing up their hands and saying, "I give up."

Steve: "Risk minimization" - like, okay, I'm just not going to go there, I'm not going to click on that link, I'm afraid, I don't know what's going to happen. "And decision avoidance" - it's like, god, no, I just don't want to have to decide, don't make me. "All characteristics of security fatigue. The authors found that the security fatigue users experience contributes to their cost-benefit analyses in how to incorporate security practices, and reinforces their ideas of lack of benefit for following security advice."

Leo: Yeah. It's not surprising.

Steve: No. No, I'm sure all of our listeners have friends and family who are just like - who just glaze over.

Leo: We all recognize this, yeah, absolutely, yeah.

Steve: Yeah. It is, it's a pain in the butt the way things are right now. So anyway, I just thought it was interesting - security fatigue, resignation, loss of control, fatalism. And it's sad, you know, this is what we've done. As if http://www. wasn't bad enough. On top of it all it's like, in the news, 500 million accounts lost by Yahoo and so forth. Oh.

Anyway, two bits of news about our two favorite browsers. Someone tweeted me, actually several people did because they know that I'm a heavy tab user, that Firefox is experimenting with native tabs, native side tabs. It's always, of course, had them across the top. There's something called Firefox Test Pilot which is where the Mozilla team stages experiments. Something called Tab Center is one that is being looked at right now which uses native side tabs. I use an add-on called Tree Style Tab, which also, as its name implies, allows a hierarchy of tabs so that I can group tabs and then close them up so that I don't have to look at actually my 222 currently open tabs. I can see a subset of them.

Also in this Test Pilot at the moment is something called Page Shot, where screenshots are built right into the browser. And once again I've solved that problem. I use a really cool Firefox add-on called Screengrab. And what's slick about it is it grabs the entire web page, not only the visible portion. And you can tell it which you want. But normally I'm wanting to grab the entire web page. And if it's a multipage document or a big scroll-y page, I just, you know, one grab, and I've got the whole thing. Which is really handy.

And so they're looking at - and I assume that Page Shot would have the same sort of features. So it may be that they're looking at some of their more popular add-ons and considering moving those features native. There's something called Min Vid which would be native support for a sticky video window which would remain visible and playing while users browsed other tabs and used the browser. They're looking at tracking protection being built in; something called Activity Stream, which they say makes it easier for you to go back and find things that you were doing. This was an interesting one. The item is called No More 404s, where it auto links the browser to the Wayback Machine. So if you have a dead link that was once live...

Leo: Oh, I like that.

Steve: Isn't that cool? Yeah.

Leo: You get the historical page.

Steve: Exactly. Rather than just saying that link is no longer valid, the browser itself will go tap into the Internet archive and say...

Leo: Love that.

Steve: ...this isn't the current one, but this is the last one that the archive took a snapshot of. Which I think is very clever. And then Universal Search, which is their, well, they have a name for their - where you're typing it in, and it's dynamically showing you things. And so they're...

Leo: Autocompletion or...

Steve: Yeah, that kind of thing. They call it the happy bar or something. I don't remember their name for it. Anyway, so that's good news on the Firefox side, things moving forward. Chrome has made an announcement about memory which will be a comfort to those who want to use Chrome, but are struggling with the bloat, which many people are talking about. Oh, that sounds like Trump. The bloat of memory consumption.

Leo: I like how you snuck that in there.

Steve: So this is due in December. It'll be Chrome 55. And what they've done is they've done a major rework of their V8 JavaScript engine for both desktop and mobile. So users of Chrome on even mobile platforms should see a significantly reduced memory footprint, by as much as 40% on sites such as Reddit, Twitter, and The New York Times, which are large heavy sites. So a couple more months, with Chrome 55, I think users will notice a dramatic reduction in memory consumption, which is great.

So we've been tracking, and I think you and I must have talked about it, Leo, before your vacation because I know that Robert and I did also, the ongoing woes of WoSign. So today - it's in the past because it was in London. So in London time Tuesday, today, representatives of WoSign, StartCom, and WoSign's parent, because WoSign is actually a subsidiary of - I guess I pronounce this Qihoo, Q-I-H-O-O, 360? I'm not good with those Chinese names.

But anyway, they all met in London to discuss this problem because we've talked about and we've shared on the podcast, I think this is a really nice, a perfect example of the problem that a web browser has, due to the power of their decisions about which certificate authorities to support and which not to. I mean, it is truly life and death for a company whether their root keys are available to the browsers that users use because, if that's not the case, none of their certificates are going to be trusted.

So what Mozilla was proposing and is apparently continuing to do is they would honor previously signed certificates, but essentially give WoSign a timeout for a year, saying we are going to no longer honor certificates signed after a certain date, and we will revisit this a year from now. So this, of course, got WoSign's attention and the parent's

attention.

And remember that, just to remind our users because we've been talking about this now from time to time, we've discussed several of the problems that the WoSign website certificate issuing system had. One was that, remember, users didn't have to use port 80 and 443. They were able to name the port where they wanted to provide authentication, which in retrospect was just crazy because it's only the ports below 1024 which the kernel has control over. And running processes are able to open ports above 1024. So that would mean that a non-privileged process running on a server within a domain could open up its own port where it runs a web server and then obtain a certificate for the privileged real domain that it's running on. Bad idea.

And then there was one where, if you obtained a certificate for a subdomain, their system allowed you to also get the parent domain. So charliebrown.github.com, you could demonstrate ownership of that because that was your subdomain on GitHub. But WoSign would give you a certificate for GitHub.com, which everyone thought was a bad idea, too.

Leo: Yes.

Steve: So, yeah, not good. So what's happened is - and it was a fascinating read. I have the link in the show notes for anyone who's interested. They finally generated a full disclosure document. And it's revealing. First of all, it demonstrates, I mean, one of the things that Mozilla took issue with was not only did WoSign not acknowledge these problems, did not report them when they were informed of them to headquarters essentially, to the CA Browser Forum group, where the baseline requirements under which they are entitled to be trusted, they were just in violation all over the place.

Then there was a sense of "You're only telling us about things we tell you about." But you've got flaws in your system which may have been exploited, we don't know how many more times. We can't trust you. And again, if there's a weakness in the CA system, it's that it is about trust. And so the players have to be trustworthy in order for any of this to make sense. And of course the other weakness is that it's a trust anyone, that is, any WoSign certificate for any domain, unless the certificate is pinned, as Google's tend to be, will be trusted.

So reading through this document, it turns out there were even more problems. First of all, most of what they explained were just sloppy coding bugs. Sometimes it was the person who wrote the code didn't understand the rules, was their explanation. They had another one, which was news to me, we hadn't covered before, where they wrote: "This is another system bug that, when the subscriber finished the domain control validation, he/she can use a special professional method" - whatever that is - "to add other unvalidated domain to the order. Then our system issued the certificate including all domains in the order." So they're saying...

Leo: Oh.

Steve: You prove ownership of one domain.

Leo: And you get them all.

Steve: So now you get the green flag, and it says, "Then you use the special professional method."

Leo: Special professional method, exactly.

Steve: Yeah. Maybe that's the advanced link, the advanced button. And then, you know, are there any other domains you'd like to put in the certificate while you're at it? Oh.

Leo: Mm-hmm, yeah, all of them.

Steve: Yeah. Give me all of them. What have you got? Oh, goodness. And then there was - this one was just - I had to read this a couple times and then dig in a little to make sure. They were actually putting ads in their certificates.

Leo: What?

Steve: They had a buy, B-U-Y, dot wosign.com advertising link.

Leo: In the cert?

Steve: In the cert.

Leo: It's like you examine the cert, and if you'd like more of these special professional services, buy.wosign.com. Wow.

Steve: So the CEO of WoSign, how did they put it, has had his responsibilities changed.

Leo: Oh. Wow. Oh.

Steve: Yes. He's out. WoSign and StartCom are divorcing. Turns out that having one entity issuing certificates under multiple names is also against the baseline requirements. Which they didn't worry about at the time. So now they're splitting up, and they're each going to have separate CEOs and management teams. They're demerging.

Leo: It's a Chinese company?

Steve: Yeah. And anyway, so this is the next stage. I don't, I mean, the meeting was

very good. The WoSign people and StartCom and the parent, their...

Leo: Their response was basically, we didn't read the manual. But now we know, so we won't do that again.

Steve: But we really liked making money.

Leo: We didn't read the manual.

Steve: Selling these certificates.

Leo: Wow.

Steve: Yeah.

Leo: So that is kind of a flaw in the process.

Steve: Well, yeah. It's a perfect example of what can go wrong with a system based on trust. It's only trust.

Leo: Yeah. So you just apply and say I want to be a CA, and they say okay?

Steve: Yeah, I mean, you have to...

Leo: Is there some sort of vetting?

Steve: Oh, yeah, yeah. You have to jump through hoops and demonstrate that you're able, that you have the infrastructure to support the certificates and so on. The problem is that there are so many of them. And again, the dilemma is that none of the browser vendors want to be the heavies. They don't want to say, you know, we don't like the color paint on your building.

Leo: Right. Just because you're in Hong Kong and you're the Post Office doesn't mean...

Steve: Right.

Leo: Right. I understand that. They shouldn't. That's right. Speaking of which, when I was gone, did you talk about the Department of Commerce giving up, ceding

control of IANA and ICANN?

Steve: Yeah.

Leo: You did.

Steve: Yeah.

Leo: And I presume that your conclusion, I mean, it happened October 1st, and the Internet's still working. So I presume your conclusion was it's okay.

Steve: Yes.

Leo: Yes.

Steve: Yes. Essentially, I understand the argument for why give up something we don't have to. Isn't the U.S. a better manager of this than an international body? But the fact is it's already international. All of the key management and root signing and everything, it's already being done in a multinational fashion.

Leo: It was anachronistic to have the Department of Commerce. And they reasonably asked for assurances that no other government would then step in.

Steve: Right.

Leo: Right. So as long as it's not a government, and it's the stakeholders that are responsible for this, and the engineers and so forth, that's the way it should be, I think.

Steve: Well, and we did, we really drilled down into it.

Leo: Good.

Steve: And what it actually was. It was a contract which ICANN had with a division of the U.S. government to perform some of the management duties, and the contract was expiring. So it was simply that it was not going to be renewed, that is, ICANN had been subcontracting. See, because ICANN is already multinational.

Leo: Right.

Steve: It had been subcontracting a part of its job to an entity that was in the U.S. And they said, okay, we're not going to renew the contract.

Leo: You and I are old enough to remember that it was a professor at USC that ran the whole thing for years. Jon Postel did the whole thing.

Steve: Yes, and his name is all over the early RFCs, when I was writing the early Internet protocols. It's like, ah.

Leo: He's passed away since. Otherwise I'd get him on in a heartbeat for Triangulation. But it was such an ad hoc thing all around. And it's just, you know, it's just...

Steve: Well, and as we know, it has functionally scaled beyond anyone's wildest dreams. And that's the problem.

Leo: I only bring it up because of the WoSign thing, which is that it is all kind of ad hoc. And so, you know, but stuff happens.

Steve: Well, and you'll remember the podcast where, when we came on the podcast, we began the podcast, and I said, "Leo, I just looked inside my XP's" - and that's when XP was new - "my XP's CA list."

Leo: I remember this. Who's the Hong Kong Post Office?

Steve: There's 400 things in there. There used to be seven. What happened? Oh.

Leo: Yeah, yeah. That's all right. It's working. It's working. And the thing is, it's self-healing. Right? It's not like Russia could come along and say, "Okay, the Internet, it's ours. We're taking it now."

Steve: Right. Yeah. The only thing that they can do is play with traffic at their border.

Leo: With BGP and mess stuff up, yeah.

Steve: Yeah. And if they started doing that, then their broadcasts would get blocked by...

Leo: Yeah, we'd cut them off, yeah.

Steve: Yeah. I mean, so it's just - it's, I mean, essentially people, you know, individual

dictators and, well, and non-dictatorial governments would love to have control. But it's too big for them to have control.

Leo: They can't. It's like controlling the ocean. Nobody owns the ocean.

Steve: Yeah. And what's interesting is, unlike electricity, where you say, okay, well, yeah, electricity is a local delivery system, the value of the Internet is that it is global. And so if you cut yourself off from the rest of the world, it would be you that suffered.

Leo: You've got the Russia Net.

Steve: Yeah, the Russki Net.

Leo: Russki Net.

Steve: So I did just want to mention, in case there was anyone who didn't know, that there have been a couple reports which are still being researched more deeply, that the replaced Samsung Note 7 phones, or a couple, like three of them, I think, so far, may still have a problem. And so Samsung has stopped - they've halted production of the Note 7. And it may just never come back.

And reading between the lines, I'm wondering if this wasn't more of a firmware problem than a battery chemistry problem because everybody knows about the Hoverboards that were exploding. This is the same technology. We take for granted the way batteries work. But there is a huge amount of energy that is chemically stored in a battery. And what we want it to do is to come out in a nice, even flow of electron pressure from its two connections. It can come out as an all-at-once chemical flow.

Leo: Right.

Steve: And that's not what we want.

Leo: That's an exact analog to the gas tank in your car.

Steve: Yeah.

Leo: Anything that stores energy, a considerable amount of energy, is potentially explosive if it comes out too fast.

Steve: Right, right. And that, of course, is the problem, the potential problem with supercapacitors is we like them because we can charge them fast, and they have no chemical-limited cycle life. The problem is there's still enough energy to push your car - in a big supercapacitor system - to push your car for hundreds of miles. That's an

amazing amount of energy. And if instead it is somehow released all at once, you don't want to be anywhere in the neighborhood when that happens.

Leo: Yeah, a gas explosion.

Steve: So very much like a full tank of gas exploding.

Leo: There have been, let's not forget, cases where gas tanks have been not safe. The Pinto; right?

Steve: Yup.

Leo: I mean, this is just more like that. I feel bad for Samsung because I think the Note 7 is dead.

Steve: I think it is. And maybe we need, I mean, I hope there are lessons that come out of this. If there's something that can be learned about, like, what it is in the battery chemistry or the charging algorithm that is weakening the battery so that it can be prevented. That would be good. I assume that the problem with the Hoverboards was just poor quality, as absolutely low cost as possible, just pump a gazillion out of these before Christmas.

Leo: I think WoSign was making the batteries, and they just forgot to read the manual. The thing that's scary, and this was true also of the Hoverboards, is it wasn't just when the Note 7's were being charged. They could be in your pocket, and they would just spontaneously erupt. And that's really scary.

Steve: Yeah. Well, and so what could have happened is that there was an overcharging which weakened the electrolyte, essentially, and caused the problem.

Leo: Oh. So it happened during charging, but the impact was felt later.

Steve: Correct. Well, and we've also noticed how warm the batteries get when they're in use. That's not just the processor burning the juice. It's because batteries have what's known as an internal resistance. And so when you move a lot of energy across that internal resistance, what gets hot when you run current through a resistor, we call those "heaters." We all remember the old-school red coil winding heaters.

Leo: Right, it's how your toaster works, yeah.

Steve: Right. And so the battery itself is made hot when it is in use and discharging. And so that could just be the straw that broke the battery's back.

Leo: That actually makes sense because Samsung's advice is turn it off.

Steve: Right.

Leo: And apparently they are safe if they're not on. So it's not a physical hazard. Yeah, that makes a lot more sense. Although XDA Developers' site had a picture today of the box Samsung will send consumers if consumers want to turn their phone in. It's a fireproof box. It includes gloves for you to wear. And then it says in bold letters on the front, "Contains potentially explosive lithium-ion battery. Not to be put on an airplane."

Steve: How do you get to them? You put it on a boat?

Leo: By truck. Slow boat.

Steve: Wow.

Leo: So, I mean, I'm glad they're taking it seriously. They should.

Steve: You know, and what's so sad is, as you have noted, it is a very nice device. It's probable...

Leo: It's wonderful. So sad.

Steve: It's probable that the majority of them would never have a problem.

Leo: Right.

Steve: But who can take that chance?

Leo: Right. It's one in 30,000 or something like that. But that's too big a chance; right.

Steve: I got a nice little bit of errata from a Peter Brumby, who said: "Hi, Steve. Love the podcast and SpinRite. It is the first and probably only time I will be able to contribute to the podcast." Well, but he doesn't know that. He says: "In SN-580," so that was last week, "you said you submitted grc.com and www.grc.com to the HSTS preload list. However, you can only submit the base domain, grc.com. The HSTS preload list is then applied to the base domain and all its subdomains, including www. It's a very minor point, but could mess up a non-HTTPS subdomain if you weren't aware. Keep up the good work."

So Peter, thank you. What's actually happened is there's been some evolution of that. I did originally submit both because, when I did it, it was before there was a list. This was just to Google themselves. Google was the first browser to say we're going to do, very much as they'd had for themselves, we're going to make this available to third-party sites. And what they wanted was explicit domain names. Since then, it's been generalized.

And I did go poke around the list and verify what Peter said, which is true, which is now it's only the base domain. The reason given was two things. One is that the way cookies are handled, there could be some confusion, that is, with cookie security, if subdomain security was allowed to differ from the base domain security. And the second is the size of the list. So many people are wanting in that they have said, okay, look, we don't want your 27 subdomains. We don't have room on the list. So we're just going to make it all or nothing. If you want to be on the list, you've got to have your entire site secure. Which is much easier to curate anyway. So thank you for giving me the chance to update everybody, Peter.

Leo: Continue on with miscellany. We're going to talk about Yahoo in a minute.

Steve: We will. And about the problem with primes. I just did want to check in about "Westworld."

Leo: Oh.

Steve: Wow.

Leo: You know, it's hard to live up to the hype; right?

Steve: It is.

Leo: And we've been waiting for this all summer.

Steve: Yes.

Leo: And it was highly hyped. And yet I think it did.

Steve: I am enraptured. IMDB has given it a 9.2.

Leo: Wow.

Steve: Which is, like, you know, to get 9.2, everybody has to either say 9 or 10, essentially. It is visually wonderful. Anyway, it's science fiction, robots...

Leo: Ed Harris. What could go wrong; right?

Steve: What could go wrong?

Leo: Now, you said you had one quibble with it. Now, I don't want any spoilers here.

Steve: No, this is not a spoiler. And you know I won't do spoils. And anyone watching the first episode would - it began to bother me in the first one, and then it really bothered me last Sunday, the second one. And that's the one-sidedness of the fact that the humans can't get shot. And but you can shoot the robots.

Leo: Well, what kind of amusement park would it be if the customers could get hurt?

Steve: Well, it would be more interesting.

Leo: Well, it would be.

Steve: I mean...

Leo: But not for the customers.

Steve: The robots could be bad shots. Or maybe you get drugged or stunned or you lose a week of your life, you're in a coma for a week or something. I don't know. But when you see them, it's like, you know, a guy just standing there being shot from all sides and ignoring them. It's like, okay, this is a problem. So that's the only thing that bothers me is that, I mean...

Leo: That doesn't bother me at all.

Steve: Okay.

Leo: I mean, first of all, it's not central to what is clearly developing as the point of it.

Steve: Oh, well, and we should say that there are multiple plot lines going on, too.

Leo: It's really good that way. And because you have the amusement - look, we're not saying anything people don't know because it's based on that 1973 Yul Brynner movie, "Westworld," which was an amusement park where the cowboys were robots,

the humans would go there, but what happens if everything goes wrong kind of a scenario. But this is going to be, I think, a lot deeper than that, and very interesting. And because you have these amusement park plots, as well as the stories that the guests participate in, as well as the subplots, it's multilayered. I think it's going to [crosstalk].

Steve: Well, yes. And the whole exploration of what it means to be "real," unquote, that is, versus, you know, like does Dolores, is that her name, does she know that she's...

Leo: Right, what does she know?

Steve: Some of them apparently don't know.

Leo: Right.

Steve: Like they're surprised when they get shot.

Leo: Every time. Over and over.

Steve: And they think they're doing their job. So, yeah, it's rich. If anyone has access to HBO, you should not be missing this. It's two episodes in so far.

Leo: Yeah. And it's really about artificial intelligence. This is just one of many, many, many stories we're going to encounter over the next few years as we grapple as humans with this emergence of these artificially intelligent machines.

Steve: And the special effects, the visual effects, like the 3D printing, it really - it goes a long way...

Leo: Oh, isn't that cool? Yeah.

Steve: ...to sell, it sells you on the idea that we don't know when in the future this is, but it looks real. I mean, they're, like, they're convincing you that somehow - I'm just wondering, what is the power source for these things? I'm trying to reverse-engineer them. And it's like, they're not going to tell us. But anyway, I just - it's delicious.

Leo: It is. It's really fun.

Steve: And I'm less excited about "Timeless."

Leo: Oh, I haven't seen it yet.

Steve: I was hoping. It's had two episodes. The second one was last night, which I have not seen. But, I mean, I think it's going to immediately fall into being just sort of a serial, like chasing around to famous moments in time, and cops and robbers and stuff. So it's like, eh. I mean, visually it looks fun. They're developing some interesting storyline. But I'm not sure.

Leo: That's pretty much what Hollywood does with time travel is they just say, oh, it's really a way to talk about different storylines.

Steve: Right. Well, and that's what Star Trek was, too. As I mentioned once before, for my birthday one year, when I had a company with lots of employees, one of them pretended to be my agent and obtained the Star Trek scriptwriter's kit from Paramount.

Leo: Oh, neat.

Steve: And she had to say, yes, I'm Steve Gibson's agent, and he writes a column in InfoWorld, and he's interested in doing some - and I never had the time or real interest. But the point is that this was something - this was a behind-the-scenes look that I had never had before. And I did read through this writers' guide. And they made it very clear that this is not - this series - oh, and this was the Next Generation period.

So they said, you know, this is not about tractor beams and photon torpedoes and warp drives. This is about human drama, set in this period. So be careful you don't write an episode that is about technology because we have no interest in that. We want to write about people set in these situations. And so I thought it was interesting that even Star Trek was basically just, you know - well, and I think that's what the appeal was. People could, even non-sci-fi people could relate to the dilemmas that were being set up in that franchise.

Leo: I think any good science fiction has to be about people, not tractor beams.

Steve: Right. In fact...

Leo: That's why we like Peter F. Hamilton; right?

Steve: Yes. And I was just going to say, I am, as I mentioned months ago, I've decided, because the Lost Fleet series had been finished, the second block was the Beyond the Frontier. And so I thought, okay, it's been a long time since I read the first bunch. So I started from the beginning. And that was like 11 books ago. And I'm - whew. The good news - and then "A Night Without Stars" has been released, which is the sequel to the - what was the first one? I can't remember now. Oh, The Fallers.

Leo: Oh, that's out, that's right, the new Peter F. Hamilton. I know. Oh, boy.

Steve: And of course I can't start until I finish the Lost Fleet. But the Lost Fleet has turned into politics, which many people find annoying and boring. I am a junkie for politics, the human drama, the interaction, the interplay. So I'm on the last book of that. And then "A Night Without Stars." So I did want to mention to everyone that that second book, that finishes - and it's funny, I was listening to you. There's, you know, trilogy is such a nice word.

Leo: What is it? Biology. What is it?

Steve: Exactly.

Leo: Bilogy.

Steve: Yes, duopoly? No, it's...

Leo: There's no word.

Steve: There isn't.

Leo: There must be a word.

Steve: You know, you could say the sequel; but it's really not a sequel, it's the second half.

Leo: It's a two-part book.

Steve: We know what a trilogy is. There's got to be - anyway.

Leo: Good point.

Steve: Someone will know what it is.

Leo: Yup.

Steve: And I did, I had something that just meant so much to me that I wanted to share. Not about SpinRite, but about the Healthy Sleep Formula. And I have a lot of these that I've not been sharing.

Leo: Oh, I know which one you're going to share because I just read it on the page.

Steve: Yeah.

Leo: Wow.

Steve: So this is from a Janice Morse, who tweeted me, @sassyjan1209. She said: "Hi, Steve. My son listens to your podcast. I have not had a good night's sleep since 1978. I have been diagnosed with all three forms of sleep apnea. I have had eight sleep studies. The conclusion is that I never hit REM sleep. I live being exhausted at all times."

Leo: Horrible.

Steve: "I wake up exhausted and shook-up from nightmares. My son brought over the niacinamide and melatonin. The first night I only got up once, instead of the usual every 25 minutes, with a headache. But I gave it another night; and, wow, not only did I get up only once, but I woke up rested, ready to go for the day, and had a pleasant dream. Wow. Now I've had my fifth successful night in a row. And I'm not only looking forward to going to bed, my entire demeanor has changed. I am so excited. If I continue to respond to these natural remedies, this is a life changer. Thank you. I will pray that this information gets out to more sufferers. I certainly will do my best. Janice Morse."

And so I replied, and I thanked her for her tweet, and I told her how happy I was that this had worked for her. And I asked her if I could share her feedback publicly. And she wrote back: "Absolutely, use it. I have never imagined I would ever have energy and happiness. I had no idea what lack of sleep was actually costing me. I am giddy with hope." So as people who have followed this know, the only problem that we've had has been availability of the components because enough people are interested and listen to this podcast and have been curious that, every time I post links to these things, the online suppliers sell out.

The good news is the niacinamide has been back since the beginning of the month - or, wait, the beginning, yes, of the month. And the oleamide, which you asked me before, Leo, whether it was necessary. For some people it is. I need that third ingredient which, as I mentioned once, is a natural substance that's endogenous to us, which some researchers at Scripps Research discovered when they kept cats awake longer than they wanted to be.

Leo: Oh, that's scary.

Steve: Which tends to build up and made them sleepy. So it is sort of a natural make-you-sleepy stuff. But some people are just fine with only the first two, niacinamide and melatonin. I need the second. Another insomniac friend of mine needs - I mean need the third. Another insomniac friend of mine needs the third. So if you do, it's available also now. And I had a really good communication with the people who make it, and they said they will endeavor to keep it in stock.

Leo: Yeah. It's good for them. Man, this is...

Steve: I haven't pushed anybody to provide feedback. I've just - so I've been getting it sporadically. The problem is, it just hasn't been practical for people to obtain what they needed in order to experiment with it. Now it is. So I just did want - I wanted to put it back on people's radar, now that all the pieces are there. A couple months from now I'll make an explicit request for, you know, let me - I'd like to put together a little testimonials page, just things like Janice's feedback, to encourage people to give this a shot because a lot of people have a problem, I included.

Leo: I prefer not to take something on a regular basis. But it's nice to have around as a kind of less intrusive way of easing into sleep.

Steve: Yes, I would say, if you don't need it, that's great. I also felt I could totally relate to her mentioning of feeling, like, excited about going to bed. Instead of just, after a while, you just think, okay...

Leo: Instead of dread.

Steve: Yes, exactly. You just dread the hours you're going to spend wanting to be asleep, but not being asleep. And this just - that's just gone now. It's just like, okay, I'm going to go to sleep and sleep through the night. And get up, you know, a couple times, maybe, to empty a bladder, but then right back to sleep again.

And lastly, believe it or not, there is a SQRL song.

Leo: Oh, you're not real on the Internet until there's a song.

Steve: I'm not singing it. But I got a DM from a listener who is a songwriter, and he wanted to make his contribution. So he sent me the lyrics, which I posted over in the SQRL newsgroup yesterday. And so he's working on the song. And in fact, I don't even know it's a he. The handle on his Twitter didn't make that clear. But in any event, at some point we'll probably have a link to the SQRL song on the SQRL pages.

And I did have a quick SpinRite anecdote to share, also through a tweet on the 10th. Someone who - I guess that's "Enomaly," 3n0m41y?

Leo: Oh, how fun.

Steve: So I think that's Enomaly. Anyway, he said: "So yet again SpinRite saves the day for me. Had an old drive with baby pictures on it. Would not read at all. Knew that if I'm not able to recover pictures, wife will put me out with the trash."

Leo: Oh, oh, oh.

Steve: He says, parens: "(Kidding). Pulled out my trusty copy of SpinRite; and, on Level 2, it fixed the drive. Now I have all pictures back. Thanks for an awesome product. Just sent you a Yabba Dabba Doo for another copy because you saved the day." And I replied, I said, "Wow, well, thank you for your generosity. That's above and beyond." But so mostly I just thanked him for letting me share this with everybody.

Leo: And from - oh, go ahead.

Steve: And Leo, just because you missed this, in case this ever comes up, we have verified, and this was two different instances, I think it was, during your vacation, that people have used SpinRite to repair their phones.

Leo: Oh.

Steve: If the phone can be put into a mass storage mode, and in this case it was done with, not Virtual Drive. I can't think of the name of it. It's one of the VM tools.

Leo: VMware or Virtual Box you're thinking of.

Steve: Virtual Box, yes, Virtual Box.

Leo: That's the free one from Oracle, yeah.

Steve: Yes. Yeah, it was done with Virtual Box. The drive could be put into mass storage mode. SpinRite saw it, ran a Level 2 on it, and fixed the problems. The phone is working much better now than it was before. And Father Robert had that same problem on one of his phones.

Leo: Nice. That's good to know.

Steve: So, yeah.

Leo: Good little tool. From Stack Exchange comes, well, an answer to the question of what do you call a two-book series. It's not a trilogy. It's not a dilogy. Although there is a Greek word "dilogy." Duology is one proposed, but that's a neologism. Dilogy is a Greek word, but it means the use of an ambiguous or equivocal expression. So I don't think that's really quite right. There's an author, Dan Simmons, who writes diptychs, stories published in two halves. The author of this is obviously quite literate. It's on Stack Exchange. In fact, I googled, "What do you call

a two-book series?" And there's lots of answers out there.

Steve: But none of them sound very good.

Leo: None of them sound very good at all. Diptych. An epic cycle.

Steve: "Diptych" sounds like something you go to the doctor for.

Leo: Diptych doesn't sound good. A saga, perhaps? And also the author points out that "Lord of the Rings," for instance, which often appears in three volumes, isn't a trilogy. It's a novel that's divided into three parts, sometimes by the publisher, sometimes more. So a series of three may not be a trilogy.

Steve: So it might - so the idea of a trilogy, it's like three interconnected, but each individually complete, stories.

Leo: Right. Or it might be a three-part serial. You know, it doesn't - so, yeah, so we'd have to ask Peter F. Hamilton what his intent was. Diptychs, dilogies, duologies, series, cycles, and sagas is this answer from Stack Exchange. And the question on Stack Exchange was, "A series of three is a trilogy; a series of two is blank." This is apparently not an uncommon question, if you google it. By the way, just for complete...

Steve: It's funny, one of the things that I've been wondering in this era of electronic books is, if someone's creating a book series, why is it that the 11th book in the series has to go back and tell you everything of the back story that has been established in the first 10 books? Who...

Leo: Like somebody would just pick up Book 11 and start there.

Steve: Precisely my point. Who is going to, like, oh, look, I've got this at the used bookstore, and I don't know who the character is. Well, that's his fault for, like, not getting the first one. They're all available now.

Leo: I agree.

Steve: So I don't want to read, now I'm on the 11th cycle of what the situation is, and I've read it 11 times, and I'm thinking, please don't. Let's move on. Because I'm just thinking, today, why quickly try to recount what's gone on?

Leo: Well, in this day of, you know, the same thing's happening with TV shows in

this day of binge watching. The old "Previously on L.A. Law" has - it's still around. I notice "Westworld" is doing it. But "Stranger Things" did not.

Steve: No. And I would notice, though, that sometimes in a long-running series what they're consciously doing is reminding you...

Leo: Remember 18 episodes ago when the Kingslayer slit that person's throat? Well, now you'll understand why he's back.

Steve: Right.

Leo: I hate that. "Game of Thrones" does that every single time. Yeah, because you're watching "Game of Thrones," you go, who the hell was that? What's he doing here?

Steve: Yes.

Leo: Previously, on "Game of Thrones..."

Steve: And so sometimes you'll need a little bit of a conversation from three months ago. And then something, and it's like, oh. So I'm understanding that they're, like, trying to give us a little bit of, like, crib notes for, like...

Leo: Memory jog, yeah.

Steve: Yeah.

Leo: Yeah, sometimes you really need it. I think maybe it's the Netflix originals that aren't doing that so much. I'm trying to remember if "House of Cards" did it. I don't think it did.

Steve: We're in the final - we're coming into the last cycle of "Game of Thrones"; right? This is it, finally?

Leo: Thank god, yes.

Steve: I know. Exactly.

Leo: It's almost, at this point, it's like a forced march. I feel like I can't not watch it.

I came this far. But at the same time I'm not looking forward to it.

Steve: No, well, it's like how I am with Lost Fleet. I've got to finish because otherwise I'll never know. But I can't wait to switch to Peter Hamilton and get back to that because that's just joyous.

Leo: Yeah. That's really true. When you get a really long series, sometimes it's just - you're slogging at some point, just for completeness. All right. Now we continue on. Let's talk about Yahoo and prime numbers.

Steve: Yeah. Speaking of paying the expenses.

Leo: Uh-oh.

Steve: Oh, boy. You're not - okay. The last shoe to drop here, no one's going to believe. But first, just to rewind a little bit, we know that back two years ago Yahoo was aware of a massive breach.

Leo: Yeah. That pissed me off. Two years ago.

Steve: Yes. Which they did not confess, which they did not warn their users of.

Leo: Shocking. Shocking.

Steve: Which is incredibly irresponsible. So that was one of the obvious things of concern when a user is saying, should I stay with them or not? It's like, okay, what would it take for you to leave? If they will lose all of their user accounts, be aware of it, but decide, oh, that will hurt our reputation, so we're not going to tell anybody. Again, I understand why that's the case. But what we see on this podcast over and over and over is that the best thing anyone can do for their reputation, because everyone can make a mistake, is to say, this is what happened. We fixed it immediately. And we're taking responsibility for it and hope you appreciate that. So maybe that requires some sophistication. I don't know. I just - so, okay. So that's the first thing.

Then comes the news which Reuters dropped about a week ago, that Yahoo had been, for about the first half year of last year, from January through June of 2015, secretly scanning all customer emails for U.S. intelligence. They reported that Yahoo complied with an order received from the U.S. government to search all of its users' incoming emails in real time. The EFF in covering this said: "There's still much that we don't know at this point; but, if the report is accurate, it represents a new and dangerous expansion of the government's mass surveillance techniques."

The EFF wrote: "This is the first public indication that the government has compelled a U.S.-based email provider - as opposed to an Internet-backbone provider - to conduct surveillance against all its customers in real time." And then in quoting from the EFF I

interjected my own note, and that is that this, of course, is the logical outcome of the going dark problem, where now the majority of the traffic that's flowing across the Internet is strongly encrypted.

So in this new environment, the endpoints where traffic is both concentrated and decrypted is the place to which mass surveillance must now move. And our listeners will remember that this was why I'm becoming more bullish on individuals creating their own OpenVPN endpoints at home, rather than using a central service, because that's, as I said then, that's another concentration point where traffic is decrypted and massively available. So it is obviously foreseeable that there would be packet capture going on there.

But the EFF continues: "In attempting to justify its warrantless surveillance under Section 702 of the FISA Amendments Act" - including the other two programs that have been covered before, Upstream and PRISM, that we discussed extensively - "the government has claimed that these programs only 'target' foreigners outside the U.S. and thus do not implicate American citizens' constitutional rights. Here, however, the government seems to have dispensed with that dubious facade by intentionally engaging in mass surveillance of purely domestic communications involving millions of Yahoo users.

"The Reuters story explains that Yahoo had to build new capabilities to comply with the government's demands, and that that new code may have itself opened up new security vulnerabilities for Yahoo and its users." And I'll explain that in a second. "The security personnel inside Yahoo who discovered this previously unknown software described it as buggy, poorly written, and rootkit-like in nature."

And then finally the EFF concludes: "We read about new data breaches and attempts to compromise the security of Internet-connected systems on a seemingly daily basis. Yet this story is another example of how the government continues to take actions that have serious potential for collateral effects on everyday users."

So, okay. What happened? The security team in Yahoo, which by the way has the, I would call it "complimentary" nickname, but it was not regarded that way. They were known as "The Paranoids" because they were always saying to the rest of management, we need to do this, we need to do that, we need to - and so unfortunately they just got labeled "paranoid." So those guys...

Leo: Gee, I wonder who else I know might be labeled "paranoid" for their overarching concern about security. Hmm. Stamos was very good. I think he was very good, and I think they had a good security team.

Steve: Yes. I completely agree. So a small group in the security group under Stamos, Alex Stamos, discovered this. They believed it to be rootkit malware that somehow got into their systems through a malicious entry point. The discovery was escalated several levels until it reached Alex Stamos's desk, who at the time, he's now at Facebook, but he was the head of Yahoo security. He researched it within the rest of upper Yahoo management, discovered or learned, he himself learned that it had been deliberately planted, without knowledge of Yahoo's security people, into Yahoo's servers by others within Yahoo. And under orders the incident was closed, and all further pursuit of the case was immediately suppressed.

So that's the way Yahoo handled this. It is believed, as Reuters reported, to have been done due to a secret order from the U.S. government. Now, there has been some

recently enacted legislation which makes disclosure of this possible, and open speech advocates are pushing for clarity and visibility on this specific matter. So we may get a little more news about this in the future.

So, okay. So then the final piece. It's hard to believe. As of the beginning of this month, Yahoo has disabled automatic email forwarding for their service. Where it used to be, when a user tries to enable forwarding, they get the notice that forwarding is currently under development, even though it's been in place for 15 years. Existing forwarding is being honored. But no one is now allowed to establish new forwarding for email coming into their Yahoo accounts. And of course it's transparent as to why because the upshot of all of this, of the news of all of this misbehavior, is people want to leave. Well, the way you leave is you go create a new email account at another provider. And then you set up forwarding on your old email address to your new email address.

Leo: But not so fast, Mr. Gibson. I don't think we're going to let you do that.

Steve: Unbelievable.

Leo: Now, I presume I can set up a Gmail account, and I've been telling people to do this, and have Gmail fetch the mail from Yahoo mail.

Steve: Ah. I'll bet you can, yes.

Leo: Yeah. Now, what I don't know, and you might have to have a paid Yahoo Mail account to do that, in other words I don't know if they give you POP or IMAP settings unless you have a paid account. But maybe somebody in the chatroom would know that.

Steve: Yes.

Leo: Just because Yahoo won't forward it for you doesn't mean you can't fetch it.

Steve: Although, I guess...

Leo: They couldn't stop that or you wouldn't be able to - email wouldn't work.

Steve: Right. Well, exactly. So we assume that you have either POP or IMAP access to Yahoo accounts, that is, external access, so you're able to pull from Yahoo rather than just going through the web interface.

Leo: I should log in and see. Because I - problem is I have a pro account, so I don't know.

Steve: I'm sure we'll have some people tell us, maybe by the time we're through talking about this, in the chatroom.

Leo: There is no, by the way, button to delete my account, either.

Steve: No. And what some people have had to do is they've had to set up "out-of-office" because you can still do that. You can create an out-of-office autoresponder where you give your new email.

Leo: Yeah, say "I'm not here," yeah.

Steve: So people are having to do a workaround because Yahoo suddenly decided to take that longstanding service offline. There's only one possible reason. They want to thwart people from leaving. And I think that is reprehensible.

Leo: Horrible, horrible company.

Steve: And, finally, Verizon wants a $1 billion discount on its pending purchase because of all of this.

Leo: I wouldn't be surprised to, I mean, maybe if they can get it at a fire sale price they'll say, well, let's go ahead.

Steve: Could they abandon?

Leo: They can probably abandon it. You know, there's usually a clause that says, if you abandon it, it's going to cost you a hundred million or something like that. But I bet you there's a further clause that says "unless it's for cause."

Steve: Yes, an escape clause, yes.

Leo: And I'm sure this is good enough cause. I can't imagine a judge would say, I don't know, I mean, Yahoo has completely damaged its reputation. And what makes me sad is so has Marissa Mayer in this. She's completely culpable all the way through. She's the CEO.

Steve: Yes. It was her decision two years ago.

Leo: Just very disappointing. I really expected better from her.

Steve: Okay. Now, trapdoored primes. This is interesting. I will quote from the abstract

of the paper and discuss it in a little more detail. So the abstract reads: "We have completed a cryptanalysis computation which is at the same time a formidable achievement in terms of size, a 1024-bit discrete logarithm computation" - okay, now, okay, we'll stop for a second. What they're saying is they successfully solved the discrete logarithm problem for a 1024-bit key, essentially, which breaks it. I mean, it is the impossibility of doing that which is the security that we're relying on.

So this first line is like, what? "We have completed a cryptanalysis computation which is at the same time a formidable achievement in terms of size - a 1024-bit discrete logarithm computation - and a small-scale undertaking in terms of computational resources, two months of calendar time on between 2,000 and 3,000 cores." Okay, now, that's earthshaking that you could crack a 1024-bit discrete logarithm problem in two months on a university research-scale parallel processing system.

"In comparison," they continue, "the 'real' record" - and they put "real" in quotes for second, I'll explain why - "for discrete logarithm," that is, the only one that has actually ever been cracked, "is 768 bits," which was announced earlier this year, in the spring, and that that 768-bit discrete logarithm computation "required 10 times as much computational power." So then here's the kick. "To achieve this dramatically faster cryptanalysis of a much harder 1024-bit prime" - and remember, when we add bits, it doesn't go up linearly. It goes up by a power, essentially by a factor of two. So 768 is vastly weaker than 1024. And we're all relying on 1024 today.

So, "To achieve this," they wrote, "we cheated. Deliberately. We chose the prime number which defines the problem to be solved in a special way, so that the computation can be made much more efficient. However, we did this in a subtle way, so that the trapdoor we inserted cannot be detected." In other words, they demonstrated for the first time ever that it is possible to carefully choose a large 1024-bit prime in such a way that it is dramatically weakened to someone who knows how it was chosen, but that that weakness cannot be seen: a "trapdoor" in crypto parlance; a "backdoor" in popular parlance.

They wrote: "Unfortunately, for most of the prime numbers used in cryptography today" - and this is an exaggeration, like I'll explain in a second, and that's why we want to clarify this - "we have no guarantee that they have not been generated with such a trapdoor. We estimate that breaking a non-trapdoored 1024-bit prime" - that is, a solid, properly, non-maliciously created prime - "is at least 10,000 times harder than breaking our trapdoored prime was for us once we knew the trapdoor."

In other words, what they have demonstrated is the feasibility of creating a deliberately weakened large prime which they know how to break in one ten-thousandth the time that it would normally take. And that's enough of a difference to mean that somebody, a state-level actor with a massive computational facility in Utah, next to a river to cool its heels, would be able to do this.

So they continue: "Our computation raises questions about some Internet standards" - and here I completely agree with them - "that contain opaque, fixed primes. Theoretically, we know how to guarantee that primes have not been generated with a trapdoor, but most widely used primes come with no such public guarantee. A malicious party who inserted a trapdoored prime into a standard or an implementation would be able to break any communication whose security relies on one of these primes in a short amount of time."

Okay. So when GRC creates a certificate, one of the things we - and anyone who has created a CSR, a Certificate Signing Request, on their server. You run the certificate

generator. And it's kind of cool. On the various Unixes that I've used, you see little asterisks being printed out as it's working. Basically it's running a pseudorandom number generator and testing the primality of big long numbers that it's coming up with. And so it's pseudorandomly using good entropy, hopefully, generating primes. We have talked in the past about that process not being random enough. That is, if you freshly boot the machine and immediately ask it to generate a new key, it hasn't had enough time to acquire entropy. And we've actually found on the Internet instances of collisions where independent parties both came up with the same really big key, which is not good.

So, but the point is you use a pseudorandom number generator, and you just start spitting out guesses, and you test each one's primality until you find one that is prime. Then that's what you use as the private portion of your key. You choose another one. You multiply them. That gives you the public key. You send that off to be signed by the certificate authority.

Now, okay. So they're not suggesting that there's any way that that's being misused. There's the problem of, as I said, not having enough entropy, but that's a different problem. But there's another place where primes are used, and that's in, for example, the TLS handshake, where you use perfect forward secrecy with the Diffie-Hellman key exchange. The standard itself describes and specifies the prime that is being used for that. That's not something that needs to be a secret with Diffie-Hellman. The whole idea is that you're able to - both ends are able to exchange the information. The attacker can know what is being exchanged, but they still cannot perform the computation, even seeing all of the packets going back and forth.

So the beauty, I mean, the elegance of this system is that it requires no secrecy. But nobody clearly described where that prime came from. Now, this may be reminiscent to that dual elliptic curve problem we had a few years ago. Remember the RSA, it was the default pseudorandom number generator algorithm, even though it was slower than the other good ones. It was based on an elliptic curve with unknown prominence. It was just, you know, the NSA maybe, we never knew, but no one could say how it was arrived at. It was just, here's the curve you use for this dual curve random bit generator, and everyone was using it until some people said, wait a minute.

So 25 years ago the question was raised, are primes safe against a trapdoor? And the cryptographers of that era knew that there was a theoretical weakness. But 25 years ago is long in terms of computational power. Earlier in this podcast we were talking about AI and the crazy amount of power that's available, the fact that a university research team can have 3,000 cores, you know, of high-speed GPUs cranking on something. You know, 25 years ago we had PDP-8s whose lights were blinking. I mean, that thing could not do crypto. So what was a theoretical, as we've often seen, what was a theoretical concern at the time was dismissed and largely forgotten. These guys said, "Okay, let's revisit this."

So here's the concern. We have no evidence that any prime in any standard in any wide use has a trapdoor. We never really knew that the dual random bit number generator, we never knew that it was a danger, but it was a concern because we knew it could be. One of the reasons that I liked and chose Dan Bernstein's elliptic curve for the SQRL protocol is he explains exactly where the numbers he chose to use in his curve came from, how he got them, and where they were derived. And that's what we need, and that's what's missing from the existing standards.

And, for example, one of the ways that this could be solved would be to show the random seed which was used to - or show the seed which was used to drive the pseudorandom number generator, which after generating a whole bunch of non-primes arrived at a prime. The point is, nobody, if it's a good pseudorandom number generator,

there's no way to reverse-engineer the prime you arrive at back to the seed which 10,000 numbers later generates the prime. So if you show, here's the prime we want to have in the standard, and here's the seed and the pseudorandom number generator that found that prime, then there's no way to engineer a trapdoor into the result.

The other interesting way is for most of the bits of the prime not to be under anyone's control. So, for example, you take the bits of pi, or the bits of "e," and you compose your prime of all of those that you can. And then you start fiddling with them in one small area until you find a number which is prime, only by modifying a few of the bits. And again, the type of manipulation required to insert a trapdoor would be thwarted by that. So we have good means of solving this problem. We just haven't been applying them.

And so I love that this research came out. I think what this means is that, moving forward, any [audio dropout] in the future, in order to be accepted as part of a standard that involves "magic numbers," as they're often called, will have to show us, as part of it, where those numbers came from. Not just, oh, you know, and a miracle happened. And after I hung up the phone with the NSA, I realized what my prime should be. No.

So the good news is, even though a 1024-bit prime is weakened by their estimate of about 10,000 times, if we switch to 2048 bits for our crypto, which is where we pretty much are now moving forward, even a 10,000-factor reduction in strength is still not a concern because, again, remember, it's not twice as hard when you go from 1024 to 2048; it's bizarrely, incredibly exponentially. It's 2^1024 times as hard, roughly. It doesn't actually scale linearly like that. But it's vastly harder.

So one thing we could do is to now, is to sort of more strongly deprecate the use of smaller primes, knowing that just going to two k-bit primes makes us safe because - but again, the idea of a trapdoor still is a problem. Maybe, for example, there are stronger trapdoors, that is, there are entities in the world that know of other means to put trapdoors in primes, that is, to deliberately choose a prime that has some special features they're able to leverage that give them more than a 10,000-factor weakening.

So I would argue, regardless, we absolutely have to know the provenance of any magic numbers moving forward, that someone came up with it somehow. And that process of coming up with it now needs to be transparent. It has to be demonstrated that anyone else, starting with the same seeds, could have, would have resulted in finding the same prime, which demonstrates it was not subject to manipulation. So a really nice piece of research. And this is really what you want. You want this kind of a canary to say, you know, to be, like, getting a little woozy on its perch, and for you to say, oh.

Leo: Pre-death canary.

Steve: Yes, the woozy canary.

Leo: Yeah. Not dead yet, but any day now.

Steve: Yeah. So just a great piece of research. Again, so I just want our listeners to have this context because the press is going bezonkers over this. You know, like, oh, the primes all over the Internet are compromised. It's like, no. As far as we know, none of them are. But we now know that it's possible. So moving to larger ones is an immediate solution. And, moving forward, we have to know where the magic numbers came from.

Leo: As always, Steve makes everything crystal clear, and I highly recommend that you listen to each and every show so that, when somebody comes along, your mom says, "Honey, what's all this I hear about trapdoor primes," you can explain. Okay. Let's see. We're going to be back here - this is it; right? This is it. This is the conclusion of this episode.

Steve: Yes, sir. And we'll see what news the week brings. If not sufficient to engage our listeners...

Leo: To derail...

Steve: ...for two hours, then we'll do some more Q&A.

Leo: That you can do by of course following Steve on Twitter, @SGgrc. He's wide open to DMs and questions there. But you can also go to the old-style GRC.com/feedback. Actually, while you're at GRC, pick up a copy of SpinRite. It's the world's best hard drive maintenance and recovery utility, and Steve's bread and butter, keeps him in PDP-10s or 8s or whatever those are behind him. So, you know, throw the guy a Franklin.

Steve: It pays the bills and keeps me able to do everything that I like to do for everybody.

Leo: Everything else, like this, yeah. And SQRL and all the other freebies there. It's just an amazing site: GRC.com. He also has audio copies of the show and transcriptions: GRC.com. We have audio and video at our website: TWiT.tv/sn. But I think the best way to do this would be subscribe. That way it's always downloaded as soon as it's available, and you just, you know, it's a nice feeling.

Steve: It's like a TiVo season pass. I would go crazy if I didn't have - just put it in and then don't worry about it.

Leo: Then Wednesday morning, you get up, you say, what can I listen to on the way in to work? Oh, I've got a new Security Now!. Just happens automatically. And you can do that on iTunes, and Google Music has it, Stitcher, Slacker. I can go on and on. Just find anywhere that you can get podcasts. Find Security Now! and subscribe. That's the best way to do it.

Have a great week, Steve. We will be back next Tuesday at 1:30 Pacific, 4:30 Eastern, 20:30 UTC, to renew the conversation.

Steve: Yes, sir. Welcome back, and I'll talk to you next week.

Leo: Right.

Steve: Bye.