BULLRUN: Breaking SSL

Description: After catching up with the week's more interesting security news and Steve's miscellany - such as NASA working on an FTL Warp Drive! - Steve and Leo take a closer look at "BULLRUN," the NSA's code name for their encryption-cracking initiative, to speculate upon just what the NSA might be doing (and be capable of doing).

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-433.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-433-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with lots of security news, including that nuclear code that was eight zeroes. And we'll talk about something called BULLRUN, perhaps a threat to SSL security. Steve has the details next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 433, recorded December 4th, 2013: Breaking SSL.

It's time for Security Now!, the show that covers your privacy, your security, your health and wellness online. Here he is, Mr. Steve Gibson, the Explainer in Chief.

Steve Gibson: Mental and physical health.

Leo: Yeah, we do it all, from Vitamin D to honeypots and more.

Steve: Oh, and we've got a variety today. We have launching nuclear weapons with the passcode of all zeroes.

Leo: I have an email from somebody about this, somebody who actually was there.

Steve: It was a tweet frenzy all week about this. We have the annoying return of air gap-jumping malware hysteria. I wouldn't say "hysteria." That's overstating it. But the people at Fraunhofer, the people who did the MP3 encoder, the original audio psychoacoustic guys, have demonstrated a network that functions, so that's freaked people out a little bit. We have the not-so-portable car killer. I have to talk about

Amazon's Prime Air service just briefly. And speaking of moving through the air, NASA actually has an experiment for a true warp drive, Leo. There's been a breakthrough..

Leo: Wow.

Steve: ...in faster-than-light travel. I'm going to talk a little bit, briefly, about a seasonal driven bedside quest of mine, and a discovery that I want to share; two interesting Kickstarter projects; and then of course our main anchor topic is to talk about BULLRUN. We learned about BULLRUN for the first time, which is the NSA's codename for their breaking SSL project, or breaking encryption, how do we break the encryption of the Internet? And Matthew Green, who's a cryptographer we've talked about often who's at Johns Hopkins, he's been involved in an alternative to Bitcoin called Zerocoin, which I'll probably take a look at on the show before long. It's still evolving. But he pulled himself away from that because he wanted to get back to addressing sort of what he felt was a loose end, which is what could they be doing?

So he did this really neat sort of brainstorming tinfoil blog post on Monday. And I thought that it's perfect foundation for us talking about what are the things that the NSA could be doing? Because they've talked about, like, having compromised SSL. And that, of course, upset everyone. And so we're going to take a look at, from Matt's post, what's possible.

Leo: Good. Boy, there's a lot to do. Steve Gibson, Leo Laporte, and eight zeroes.

Steve: Okay. So, yeah. So the idea of this is just so bizarre. The story broke that for 20 years the U.S. was using eight zeroes as the so-called "nuclear launch codes" for the nuclear arsenal. And this came from a site called TodayIFoundOut.com. And the title of the posting was "For Nearly Two Decades the Nuclear Launch Code at all Minuteman Silos in the United States Was 00000000" was their deal.

So the background here is that we originally had no protection for, I mean, other than all of the sort of standard - these things are not out on the street corner. They're in highly secure bunkers and silos. But JFK, our U.S. President at the time, in 1962 he put out a National Security Action Memorandum 160, which required that there be some, essentially, passcode, password protection on nuclear weapons, and that they be really functional. They're a thing called a PAL, P-A-L, and that stands for Permissive Action Link. And during my research I found a really interesting, I mean a really interesting paper. You can see the link in the show notes, Leo, that www.cs.columbia.edu link, where they discuss what is known in the open community, not triple-top-secret and so forth, about this whole PAL technology, that is, this whole notion of controlling the accessibility, essentially, of a nuclear warhead's explosion.

And what's interesting about this, I think, is that, well, first of all, where the 00000000 came from was that the military commanders, specifically the people at SAC, the Strategic Air Command, were miffed at the idea of the politicians telling them how to do their job. So what we learned from the story is that shortly after the politicians oversaw the fitting of launch codes with the very first PAL technology, they reset the codes to all zeroes. So it is apparently true that, in fact, for two decades this extra interlock - I mean, this is not the only way of, obviously, getting to arming and engaging these bombs.

Leo: No, in fact, according to what I've read, it doesn't in fact engage the warhead. It merely allows the rocket to be launched. And there were separate codes for the warhead. In fact, we got an email from a guy who was running a silo. His name is Joseph. I probably shouldn't give his last name. I think it's okay. He says, "Your story about the nuclear codes is ridiculous. I was a launch crew member on Titan 2 missiles in Kansas." Now, he said, "Minuteman may be different, but our codes weren't numbers at all, they were letters." And I think the confusion, I've asked him to clarify, is that the letters were used to arm the warheads. The numbers were used merely to launch the rockets. He says, "As you can see in the photo, and also in the manual" - he sent me some pictures which I'm going to show you.

"Also there is my own set of launch keys I got from my site after deactivation. I've included the launch checklist and the entry for the code add in the BVL, the butterfly valve lock." And he also has a picture of himself at the nine-megaton warhead that he was monitoring. Let me just show you some of these images. And I'll send these along to you so you can read them. Because it's been deactivated, these are no longer - there's the butterfly valve lock. And you see, I think you can see where you would insert the key. I'm not sure. Here's where you'd enter the code. And this is the manual. Here he is in front of his, what did he say, nine-megaton warhead.

Steve: Yikes.

Leo: And there's the two keys. Remember, there are two keys. Each member of the combat crew, there's two of them, has to have a key to turn, and they must turn them simultaneously, and one person can't. They're separated enough. So I think that's pretty cool. But he disputes it. But I think we're talking about different things, is what I think.

Steve: Yes.

Leo: I've asked him to clarify.

Steve: That's just what I was going to say because I'm absolutely certain that what this was, was actually about the bomb being able to be detonated.

Leo: Really. All right.

Steve: Yes. In the documents they show schematics and talk about specifically how the systems were designed. And what intrigues me is this is a classic security problem. And in fact it was concern about our warheads being on foreign soil, I mean, like deliberately on foreign soil, where we - it wasn't...

Leo: Well, it started with the Cyprus problem in the early '60s, '62, where the fear was that NATO missiles would be used by the combatants against each other, even though they're both in NATO.

Steve: Yes. And so this was meant to be an interlock that would prevent the nuclear bomb from detonating.

Leo: Now, this story comes from 2004, as you know. So this is - a lot of people will say, and you'll get emails saying, oh, we've known about this for years. This was originally written about in 2004.

Steve: Yeah. And I'm just addressing the fact that it was the most tweeted thing I saw this week.

Leo: Yeah. Oh, no. It got to be a big story, yeah.

Steve: Yeah. So but from a security standpoint, I mean, this is the kind of thing we've talked about often. And think about the dilemma. The reason the Strategic Air Command commanders set this to all zeroes was their concern that they wouldn't be able to get some random code when they absolutely needed it. And the point was this wasn't the only thing you needed. Everything you just talked about with the two-key interlocked, being twisted at the same time, the stuff we've all seen in the movies, that was all in place, too. So there was a whole series of things that had to happen. But the tension from a security standpoint is that you both absolutely never want there to be a mistake and a warhead detonate anywhere that you don't absolutely want it to. And that's in tension with the idea that, if you want it to, you absolutely want to make sure that this huge list of things that all have to happen in order for that to occur, all do.

So, I mean, there's a real problem. Again, you absolutely don't want a false positive. But when you want it to occur, you want to be sure that it's going to. So it's really a dilemma. And one of the problems was solved by simply, essentially, zeroing out one of the interlocks, actually on the physical weapon itself. In this Columbia.edu paper, they really go into - it's a fascinating paper, for anyone who wants to read it. I've got it in the show notes.

And, by the way, I do have - the show notes are already posted on GRC. I just tweeted the link a few minutes ago before we began the podcast [www.cs.columbia.edu/~smb/nsam-160/pal.html]. And I got a lot of great positive feedback from my posting of last week's notes. So I'm going to - that's what I will do from now on. I wasn't always making sure, like sometimes I had people's email addresses and things in that I would never want to disclose by mistake. So I'll make sure that they're postable from now on. But anyway, it's really...

Leo: It's a great story. And not at all - to me it made perfect sense. It's exactly what you'd expect, frankly. It was kind of a pragmatic approach.

Steve: Now, of course we had badBIOS, and we talked about that a couple weeks ago. This persistent rumor-cum-belief, whatever, that one security researcher has been plagued by this weird BIOS that affects - or, I'm sorry, this weird malware that apparently affects, is able to infect BIOSes and, he believes, jump into completely disconnected computers. And the only channel he was able to find, since it wasn't networked, there was no Bluetooth, there was no WiFi, I mean, it was "air gapped," is a term, is that he believed that a laptop was still getting infected until he physically

disconnected its microphone and speakers. And then, oh, thank goodness, it was no longer - nothing was able to get to it. And it's like, okay.

And as we talked about it at the time, it's like, yes, in theory, except there are all kinds of problems with that idea. First of all, the bandwidth is going to be low. The microphone and speaker really have narrow acoustic bandwidth ranges, so it's going to sound like a modem, and so forth. So this all kind of quieted down, and then unfortunately a story just surfaced a couple days ago about, as I was talking about at the top of the show, the Fraunhofer folks have developed a network that does this. That is, and they have an interesting diagram where they talk about - it's in Ars Technica, and our friend Dan Goodin posted this: "The topology of a covert mesh network that connects air-gapped computers to the Internet."

And the reason you need a mesh network is that, of course, the distance is going to be extremely limited, actually farther than you might think. The guys at Fraunhofer were able to operate 65 feet between two Lenovo laptops using the built-in microphone and speakers. Now, reality begins to hit here. They were able to get 20 baud, that is, 20 bits per second, which is not surprising because it is audio, and so you're going to have to have some sort of a carrier. And I guess they were able to get it out of audible range so it was technically ultrasonic, or it had to be very low ultrasonic to still allow the microphone and speaker to function.

But so, okay, yes. Technically you can use acoustics to allow two in-proximity laptops to talk to each other. To get greater distance, then, they use a mesh network, which is to say essentially it's like the Internet, where one hops to the other, hops to the other, hops to the other. So as long as there are any two that - or as long as there is any path where two are within distance, then they can all be synchronized, all be at 20 bits per second of data rate, which is a little slow in terms of contemporary networking technology. But, again, all it can be used to do is allow previously infected machines to communicate. And that's the key.

And that's why - and it's not like an infected laptop at Starbucks is going to be able to reach out and propagate itself to all the other laptops within 65 feet of it acoustically because a microphone can't take over a laptop unless there's already an infectious agent, something malicious in the laptop listening on the microphone for instructions. So, yes, it could be used for persistent communications, stealthful, low-baud, persistent communications. But not to just reach out and take something over. As far as we know. As far as we know there isn't any way that just whistling to a microphone can cause a buffer overrun and install code. And, if so, it'd be prohibitively slow.

Leo: Interesting stuff. Nothing to worry about. Nothing to fear.

Steve: Nothing to worry about. So, yes, now we have some sense of - if anything, it gives us a sense of scale for what can be done in an acoustic network, if that were actually what was going on in this badBIOS situation.

The friend of the show Simon Zerafa, in Wales, sent me a link. He called it the "portable car killer," of course, playing off the well-known portable dog killer experience in my youth and episode that we did by that name. However, this is not so portable. It is 772 pounds. But it is a vehicle. It is, I mean, it is functioning. There's a company called E2V has developed this. It is a nonlethal device, unless maybe you have a pacemaker, in which case maybe it's not so nonlethal. But it can shut down a vehicle at 165 feet, at 50 meters. So basically it is a mobile EMP, electromagnetic pulse, transmitter.

What it does is essentially it takes advantage of all of the technology that we have in contemporary vehicles. If you had some old Rambler from the '60s, such an old-school car would just ignore the electromagnetic pulse. But the guys who develop this note that, in the frequency band they're using, which they call the L and S band, a contemporary car's wiring loom has runs of wire of about a meter, which make them a perfect antenna. And so they're able to generate a signal from 165 feet away, which at that distance will essentially scramble the engine's management technology and cause the car to stop. And they're claiming that, after their press release of this, they have had interest from 15 different countries; and, I mean, like dramatic interest in, okay, we want one. Where do we send our money?

And so far they've installed this 772-pound thing in a Nissan Navara and a Toyota Land Cruiser, just to make their demonstrations feasible. I have a picture of it and links in the show notes, if anyone's curious. But nothing most of us have to worry about. And I have to say that it would unfortunately knock Amazon's Prime Air delivery drone right out of the sky. It's pretty clear that it would scramble its avionics, too.

Leo: And you don't really want something that could be knocked out of the sky.

Steve: No, Leo. I mean, actually what I loved was Paul Thurrott's tweet. Paul tweeted, I think this was just yesterday, he said, quote: "The sheer amount of free PR that Amazon's CEO Jeff Bezos got for his," as Paul said, "BS 'drone delivery system' is awe-inspiring. Media, you just got played."

Leo: Yeah, yeah.

Steve: And, I mean, and I'm sure everyone listening has heard about this. If not, just Google "Amazon Prime Air," and you'll see it. And I don't know what to make of it because you can't have this thing with eight exposed high-speed rotors spinning. I mean, it would shred the family dog that would definitely go chasing it and barking at it and wouldn't see the rotors spinning and would, I mean, there would be fur flying. It's just crazy. I don't know what they're thinking.

Leo: There are lots of, well, and I think even Jeff was pretty clear that this was a long-term R&D project, not something they'd definitely like to do. And it would save them a lot of money. You know, it has been used in Haiti. Something similar has been used in Haiti to deliver medical supplies.

Steve: For delivering. Oh, interesting.

Leo: So it's not out of the...

Steve: Oh, Leo, it's technically - I think it's clearly feasible. We have the battery efficiency now. We have GPS that is ubiquitous. The cost is low enough. No, I mean, it's absolutely something you could do. And wouldn't it be fun to just look up and see these little autonomous bots. It's unfortunate that they got labeled "drones." I kept seeing this

word "drone." I thought, oh, goodness, that already has so many negative connotations associated with it that you really don't want that to happen.

Leo: One could come up with a lot of issues. But it has been used in the past. There's an article in the MIT, what do they call it, the - MIT's Technology...

Steve: Yes, Technology Journal.

Leo: Yeah, about it and about how it's been used in the past. And it's not completely infeasible. There's just some big problems that would have to be solved. But it's, yeah, I mean, the fact that Charlie Rose said, "Oh, my goodness," and completely uncritically, "Wow."

Steve: No, I mean, no, Paul's right. I mean, the media was saturated with this.

Leo: Jumped on it. It's a great story.

Steve: It was a fabulous video that they produced where it was basically a few moments in the life of our next-generation drone delivery system. It was like, oh, my god.

Leo: I enjoyed it. And you know it was a great interview and, you know. What the hell.

Steve: Now, I hate to follow that with this next story because - okay. And first when I saw this I thought, okay, is it April 1st? No, it's not April 1st. Okay, how do I explain this? So this is NASA actually has a project. And I have a link to - and if the PDF weren't being hosted by NASA.gov, again, I wouldn't believe it. But it is. They actually have - now, okay. That's just the - I think that's the...

Leo: It's an artist's rendering.

Steve: No, that's a Vulcan something or other, the very first image is, yes, and an artist's rendering from the Star Trek world. But that one you're showing now is actually what this thing might look like. Okay. So first of all, for the listeners who aren't seeing these pictures, the article begins - this is in io9.com, by the way. "A few months ago, physicist Harold White" - who's the author of the PDF at NASA, at JPL, or some propulsion organization. I'll get to that in a second. He "stunned the aeronautics world when he announced that he and his team at NASA had begun work on the development of a faster-than-light, i.e., FTL, warp drive."

Now, okay. We have to just pause and give kudos here to Gene Roddenberry because, I mean, we're calling this a warp drive. It actually works by warping space-time. And we had that in Star Trek in the '60s, thanks to Gene Roddenberry. I mean, that was the technology. That's how the Enterprise moved at faster than light. So, incredible.

So "[Harold White's] proposed design, an ingenious re-imagining of" a drive known as the, and I'm going to kind of mangle the name, I'm afraid, Alcubierre. He's a Mexican theoretical physicist, Miguel Alcubierre. I was trying to pronounce it before the podcast, but now I forgot. "[The] Alcubierre Drive may eventually result in an engine that can transport a spacecraft to the nearest star" - okay, Leo, to the nearest star - "in a matter of weeks."

Leo: Well, we need that because otherwise it takes too long.

Steve: Exactly, "all without violating Einstein's law of relativity." And I've said on the podcast that one of the problems I think that NASA has and that the whole space process has today is that we've been spoiled by Star Trek and by all of the sci-fi movies. No one is able to generate much enthusiasm about a slow wagon train to Mars. It's like, who cares? I mean, no one is going to fund that. You're just not. But oh, my god, going to Alpha Centauri in a couple weeks? That's a game changer. Now we're talking.

So anyway, so about Miguel. He's got a Wikipedia page. And back in '94, okay, so 20 years ago - Miguel's about 48 or 49 now, so he was at the prime of his physics inventing age back then, 20 years ago - he published a paper in the Classical & Quantum Gravity journal, so serious...

Leo: That's the right place for it, I think.

Steve: Yes, theoretical physics. So Wikipedia says Alcubierre is best known for the proposal of "The Warp Drive: Hyper-fast travel within general relativity." And that's the key because the problem, of course, is acceleration. You can't accelerate very quickly because we have no way yet of suspending inertia, or humans are turned into goo, and that's not good. So anyway, continuing with Wikipedia, "which appeared in the science journal Classical & Quantum Gravity. In this he describes the Alcubierre drive, a theoretical means of traveling faster than light that does not violate the physical principle that nothing can locally travel faster than light. In this paper he constructed a model that might transport a volume of flat space inside a bubble of curved space." So Leo, we have a warp bubble. I mean, again, we have lots of science fiction about this, but now we actually have theory.

"This bubble, named as hyper-relativistic local-dynamic space, is driven forward by a local expansion of space-time behind it, and an opposite contraction in front of it, so that theoretically a spaceship in the middle would be placed in motion by forces generated in the change made in space-time." Okay, now, this is 20 years ago, and nothing happened much because the theoretical amount of energy that is unfortunately required to warp space-time is quite literally astronomical. Space-time, it turns out, is very stiff, and actively resists being warped.

Leo: Dammit.

Steve: I know. It's been a big problem. But what happened was...

Leo: We're so close.

Steve: And we have such great pictures. Unfortunately, we have no way of producing that much energy. So Dr. Harold White at NASA's Johnson Space Center was putting together a presentation recently where he was going to talk about this problem. And his PDF is titled "Warp Field Mechanics 101." And if we were able to build this, this would transport a spacecraft to Alpha Centauri in two weeks.

Leo: Perfect. Just right.

Steve: Even though the system, even though Alpha Centauri's system is 4.3 light-years away. So we're talking about 4.3 light-years in two weeks. And oh, my god, I mean, now we're talking.

Leo: Now we're cooking with gas.

Steve: Yeah, exactly. If we could generate the insane amount of energy required. Well, and that's the breakthrough. What happened is, while preparing this report, he did a sensitivity analysis of the equations, and he believes he found a way of dramatically reducing the amount of energy required. And so in his paper it says, "It takes advantage of a quirk in the cosmological code that allows for the expansion and contraction of space-time and could allow for hyper-fast travel between interstellar destinations. Essentially, the empty space behind a starship would be made to expand rapidly, pushing the craft in a forward direction. The passengers would perceive it as movement, despite the complete lack of acceleration."

And he said, "In terms of the engine's mechanics, a spheroid object would be placed between two regions of space-time, one expanding and one contracting. A 'warp bubble' would then be generated that moves space-time around the object, effectively repositioning it. The end result is faster-than-light travel without the spheroid, or spacecraft, having to move with respect to its own local frame of reference."

So essentially you create a bubble. This thing is in the middle. And you sort of rotate the bubble. And the craft moves through, essentially is pushed out of our normal space-time constraints and is then able to travel without inertia being a problem and without even something pesky as lightspeed, the speed of light, being a problem, and just zip wherever it wants to go.

Leo: Zip.

Steve: Zip, yeah. So anyway, we'll keep our eye on this. But, I mean, it's real.

Leo: We'll be around for a while; right?

Steve: They're now doing - well, no. But the energy problem just stopped everyone cold.

I mean, it was absolutely - it required an absolutely infeasible amount of energy. And it's been reduced by billions of orders of magnitude. I mean, it's, like, been reduced - if this new understanding is correct, and what they're actually doing now is building an interferometer to test the warp bubble theory because they now believe, I mean, other people have looked at it, and NASA has looked, I mean, NASA's guys have said, wow, this could work.

Leo: I love it.

Steve: And so they're going to start with an interferometer in order to see whether they can actually create a warp bubble with this radically lower level of energy input required. And I haven't had a chance to read the PDF. I just found it this morning. I tweeted it. It's already in my Twitter feed. It's in the show notes. So, yeah. Wow. This would change everything.

Leo: In our lifetime, you think? Nah.

Steve: Oh, yeah.

Leo: Really? Think so?

Steve: Oh, Leo. I was looking at this, you know, there was all of this JFK's 50-year event in the last week. I remember where I was on that Saturday morning.

Leo: Who doesn't? Anybody alive does, of course.

Steve: Yes. And I didn't really understand what was going on. But I knew from the expression on my dad's face that something really bad had happened. We'd been out sailing in the Bay, in San Francisco Bay, all day. And so we were coming back, we came back to the dock in Marin County, and my sister and I were hosing off the sailboat. Dad had gone to the shore. And when he came back, someone on shore had - so we'd been out of touch with the news as a consequence of being sailing. And Dad found out when he was ashore and walked back to the boat. I mean, I'm just like, "Dad" - or Daddy probably at the time, I think I was eight, it would have been 50 years ago - "what happened?" Anyway, so the point is that seeing the film from then, look how far we've come in 50 years, Leo. I think we forget how incredibly rapidly technology moves. I mean, that was not long ago. That was 50 years. And...

Leo: But Steve, I know you're bullish about your Vitamin D and your low carb; but I don't think we're going to last another 50 years. You think?

Steve: Well, oh, my goodness, yeah. I have a T-shirt that says "Future Centenarian." Yeah.

Leo: Right on, Daddy-o.

Steve: Absolutely.

Leo: I don't see myself getting to 107. But maybe.

Steve: It's not going to take 50 years. It's not going to take 50 years. Look at the Large Hadron Collider. I mean, that thing is science fiction. Have you seen the pictures of that? It's like, we're building something like that. We can easily build a starship. All we need is the warp bubble. And apparently that's just down the street now.

Leo: I say start building the starship. And then, if you build it, maybe the bubble...

Steve: Give it a big - give it a nice big engine room, and we'll figure out what to put in there.

Leo: We'll figure out the rest. Yeah, exactly.

Steve: Exactly. And we have Elon Musk, you know, so he's not going to let anything stop him.

Leo: We've got Google.

Steve: We do.

Leo: Yeah. They seem to have a vision for the future, moon shots.

Steve: No, this is exciting stuff, Leo. I mean, this is - what's exciting is that now we can actually go somewhere in a reasonable amount of time.

Leo: Right.

Steve: And if we can do it in a couple weeks, then life support problems are solved, inertial problems. We don't have to freeze people. We don't have to, I mean, we don't need all this other stuff we don't have. All we need is the warp bubble, and apparently we're going to have that soon.

Leo: The only thing I would raise at all is that, if we could figure this out, not being the most advanced race ever, why haven't others visited us?

Steve: Actually, I'm reading a rather daunting sci-fi series at the moment.

Leo: We may be the only ones. Is that what you...

Steve: Yeah.

Leo: Is that what you were going to say? I find that hard to believe, as well, when you have as many planets, potential planets as you must have, that no life has ever developed on any of them. Seems unlikely.

Steve: Well, for one thing, we really are way out on the fringes. We're not, I mean, we're on the wrong side of the tracks, galactically speaking. And so there may be a lot more going on somewhere.

Leo: Yeah, but if there's a warp drive, we're not as out-of-town as we thought. We might be closer. It's the Fermi paradox; right? I mean, it's the classic paradox.

Steve: Yeah. Alistair Reynolds has a series called the Revelation Space series. It's rather dark science fiction, and I'm reticent to recommend it. He's also not as good a storyteller as Peter Hamilton, yet his books are just as big. And so I find myself sort of dragging myself through this, like, okay. I mean, I have to find out where this goes now. But, boy, it's not nearly as delightful. I'm on the third of the trilogy. And he answers this. And I can't say much more about that because I'd be spoiling it. But he posits in this series an entirely feasible, well, within his universe, explanation, which is interesting. It's not more true than the warp drive, but yeah. I mean, I don't care if nobody else has figured this out. But you're right. If it's this simple, why isn't everybody doing it and visiting each other?

Leo: Why aren't we getting more looky-loos?

Steve: More traffic, yeah.

Leo: You've got some nice real estate here. How much you want for the planet? You know? Moving along.

Steve: So I'm getting a huge amount of positive feedback from my Leo-precipitated change in the way I use Twitter.

Leo: Oh, good. I was feeling bad about it last week.

Steve: No, it's okay. Everyone is liking the fact that they can see my replies. I don't think anyone knew I was replying to everyone because I was using DMs. The problem I mentioned has already been fixed with that really neat filter, or that chronology of my

timeline. Remember that bit.ly/sggrc, all lowercase, took us over to Simon Paarlberg's page, where he was monitoring the feed in real-time and tying those posts to Security Now! episodes because I often tweet links as I'm preparing for the show, just so people have access to them. And I've always been saying, oh, I just tweeted this. Go check out my Twitter timeline, you can find the link. The problem was, since I was now using @replies, my timeline was hugely cluttered because I've been so active with that. Anyway, Simon immediately fixed that, and so now we're back. So bit.ly/sggrc is cleaned up again, and anyone can use that to immediately find only my broadcasts to my Twitter community. Which, by the way, topped 40,000 the other day.

Leo: Congratulations.

Steve: So I'm not where you are, Leo, but I'm...

Leo: But you say good stuff. I say nothing. So that makes you valuable.

Steve: Well, I also tweeted about this. Now, okay. This is a complete diversion, but bear with me for a second. Last year the radio station I use on my bedside clock radio switched to Christmas music the day after Thanksgiving and was playing Christmas music for a month. And I thought, no. I cannot have Christmas music for a month. I'm not a Grinch, but there actually is some Grinch song which is really nauseating, and I kept hearing "Mr. Grinch" and all this. I thought, no, I just can't put up with that.

So as the holidays were approaching, I decided to get preemptive here and needed to replace my wonderful clock radio with something. I do subscribe to XM Radio. I have a lifetime subscription to Sirius XM in my house. So I was thinking, okay, that would be a possibility. And of course we've got Pandora, we know about Pandora and so forth. But what about the device? And what I finally realized is I had an unused iOS device. I had an iPhone 4. And of course I've moved to the 5. And many people, it occurred to me, have previous versions of iPhones or older iPads that they may not be using.

What I found, and the reason I bring it up, is that it's just an incredible bargain, is an amazing little dock for an iPhone or iPad, which is also a speaker and a charging station. What I like about it is that it was - it's been discontinued, so there's existing inventory that Amazon has. It used to be $80. It was originally $79.99. Now it's $19.99 or $19.95. I've got a link in the show notes, if anyone's interested. It's the iLuv, i-L-u-v, iMM190. They call it the App Station Alarm Clock Stereo Speaker. Anyway, it's, for the price, it's terrific.

And then I went on a search for the right app to use on my old iOS device. And believe it or not, I can't find one. I've settled on something called The Clocks in the iTunes Store, which is very close. And I've written to the guys, or the guy who is the author and suggested that this would be perfect if he only made the following changes, and he said that several of them are on the way of the changes that I suggested. So if anyone's interested, they have an old iOS device for a really amazing price. And I forgot to mention that it sounds amazing. It's a bass-ported little speaker box that sounds fabulous. So it allows you to repurpose something that you may no longer be using for a very good price. So just a little hint for the holidays from me.

And finally, I just wanted to bring to our listeners' attention two interesting Kickstarter projects. Obviously, there's a huge interest in coffee. There is something on Kickstarter

called the Temperfect Mug, T-e-m-p-e-r-f-e-c-t Mug. The interesting thing about it is that the guy who designed it, he's got a lot of experience with making mugs. He's been making them for years and years. He notes that coffee starts out being too hot to drink, and you then need to wait for it to cool. And in a thermally insulated mug, any traditional thermally insulated mug, there's a certain taper to the rate at which it cools. And if you scroll down, Leo, you'll see way down he shows the temperature versus timelines of his solution for...

Leo: The guy is obsessive.

Steve: Oh, no, he really is. And there are some beautiful - you can spend $280 on one of these things that are like some amazing titanium oxide coating that you can get on one. Yup, there it is. Anyway, so the point is that a typical thermal mug slows down the rate at which the coffee gets cool. But you have to wait for it to get down to drinkable temperature, and then it drops out of that ideal zone pretty quickly. So what he's done is he puts something with a great deal of thermal inertia directly around the coffee-containing area. And then that's vacuum insulated from the outside. So the point is, when you pour hot coffee in, the temperature of the coffee immediately drops to the sweet spot of drinking temperature because that high thermal inertia, and I don't remember what it is that he's wrapped around there, but that takes the heat immediately out of the coffee, bringing it down to drinking temperature. But then it holds it there for, like, an hour and a half. So a really interesting idea. And I thought I would let our listeners know.

Leo: I like it.

Steve: Yeah, I do, too. I do, too, very much.

Leo: Seems a little bit like a perpetual motion machine or something. But if it works...

Steve: No, but, see, that's just it. The way it works is it rapidly drops the temperature of the coffee by heating up the liner that is immediately adjacent to the coffee. And then that liner holds the heat and keeps the coffee warm. So it drops it down. Instead of, like, waiting for half an hour for it to cool off enough, I have these - I really like - you keep seeing this thing. This is a Contigo, C-o-n-t-i-g-o, which I absolutely love. They're at GoContigo.com. And this thing will keep my coffee warm for a couple hours. But I transfer it into another cup because I can't drink it out of here. It's too hot to drink. And so transferring it allows it to get cool. And then the thermos keeps it hot itself. So it's a standard aluminum vacuum thermos.

Leo: So for 40 bucks you can get the regular mug. For 160 bucks you get the black oxide. And if you want titanium, it's $280. Which one did you order?

Steve: I did go the titanium one.

Leo: Of course you did. The flat blue-black ultra-hard coated mug with an uncoated droplet logo.

Steve: Yup.

Leo: Hmm. Now, they're saying summer of next year before you get this thing.

Steve: I am not in a hurry. I'm, you know, I've got other stuff.

Leo: Actually, I am so tempted.

Steve: I'm convinced that the guy's got his physics right. The physics makes sense. He's clearly a perfectionist. He's got pictures, I mean, this does not look to me like one of these things that'll never happen on Kickstarter. It looks like a deal. And for 40 bucks you still get the same performance for a somewhat less cosmetically over-the-top mug. I love the idea of it immediately dropping it down to drinking temperature and then holding it.

Leo: The Titania doesn't perform any better. It just - it's a look.

Steve: No, it's pure - yeah. It's pure...

Leo: Problem is it's either orange, pink, or blue, if you don't spend the money. Oh, I'm going to get the orange one, what the hell. You can have the fancy one.

Steve: Yeah. I'll show it to you next summer.

Leo: Yeah. Bring it with your coffee setup.

Steve: Okay. Last thing. Wackiness. But again, intriguing. I just wanted to make sure our listeners knew. A really interesting smartphone, Bluetooth Low Energy, controlled paper airplane.

Leo: [Laughing] Okay.

Steve: Now, what's so cool about this is that what - and it's just $30.

Leo: By the way, they've raised half a million almost, so they're doing all right.

Steve: I was just going to say, the goal was $50,000. They have $434,000 pledged because - and in my notes here I said, "Just the guts, ma'am," because all you're getting is this cute little armature, essentially, a little cockpit in front where the Bluetooth radio and batteries are, and a little rod that runs to the back with a rudder and a propeller. And so you get...

Leo: You supply the plane.

Steve: Yes. You supply the plane that you attach this to. And so it's like build your own plane system. I just love it. And it's $30. So no wonder that they're going crazy, pledge-wise. I didn't want our listeners to miss out on it. So I think you could, let's see, you would Google "smartphone-controlled paper airplane," that's in the URL of Kickstarter, smartphone-controlled paper airplane. Oh, that's the other thing. You twist your smartphone back and forth, you use the inertial sensor of your smartphone in order to steer your paper airplane.

Leo: It'd be fun to enter your paper airplane into some contests and let people see this thing.

Steve: Yeah.

Leo: It really flies. They have a video of one, a prototype, and it's pretty impressive.

Steve: Yeah. I don't know what the timing is. It'd be so great if it were available as a Christmas present for our listeners to give their sons and daughters.

Leo: Yeah, yeah. Price is right.

Steve: This is a tremendous, neat little - oh, it is. A great little concept. I love the idea of this is - we're only going to give you and sell you the part you can't make. And anyone can fold paper. And so, have at it. Create, I mean, like it brings paper airplanes back to life.

Leo: Remember that? You probably - we're of the same vintage. You probably had that book. Remember they...

Steve: Yep. I know the book.

Leo: Everybody's nodding. Yeah, we know that book. That was very popular in our youth.

Steve: It was, yeah.

Leo: I wonder if they still - I should look, I'm going to look on the Amazon.

Steve: Oh, it's got to be around.

Leo: You think? It's not out of print?

Steve: So speaking of bringing things back to life, I got a really nice note from a Caleb Allen that I wanted to share with our listeners. He's in Turlock, California, a listener himself. And he said, "SpinRite helping elementary schoolchildren read." And this is something I hadn't seen before. He said, "Dear Steve. I work at a small, poor, elementary school district in California's Central Valley." That's where Turlock is, sure enough.

And he said, "About two years ago I convinced my boss to purchase a SpinRite site license" - I'm sorry - "to purchase a site license for SpinRite 6 to use in our shop after I heard about it on Security Now!. Last week," and this he sent to me, by the way, on Halloween, on October 31st, so just about a month ago. "Last week one of our librarians told me of a problem where her library kiosk terminals were taking five to seven minutes to log on. It was so bad that most schoolchildren, many of whom were only given five to seven minutes to run to the library at all, just abandoned the PCs. This was a huge problem because these PCs are how the students look up book titles in the subjects they're interested in."

So I guess, back when I was in high school, we had a card catalog. But obviously that's all gone online now, as would imagine. So he said, "After an hour or so poking around the multipoint server install," he said in parens, "(the kiosks are terminals all running off one host machine)," he said, "I took it back to the shop for further work. The message stated that the user profile service was busy. I scoured the Microsoft network and Internet forums, trying to find a solution to the problem. Finally, in desperation, on the off chance that the problem might be hard-drive related, I ran SpinRite on Level 4 over the weekend. On Monday morning I came in and tested the PC, and the login popped right up and took me to the desktop. I returned it to the library, and now four students at a time are able to do research and look up books, thanks to SpinRite.

"As a poor district, we've used SpinRite for a whole host of problems. It's gotten to the point where, in most cases, we just run SpinRite on Level 2 before trying anything else. It's helped us keep our existing equipment running in the sometimes chaotic environment of elementary school classrooms. Kids are rough on the equipment, and our hard drives take a beating. With SpinRite's help we've decreased downtime, kept older equipment running, and recovered vital files such as grades, parent reports, and special education evaluations. I don't know if a student in our district will someday be inspired because they had access to a PC we've gotten running again with SpinRite, but I'd like to think so. Thank you so much for a great product. Caleb."

Leo: Tremendous.

Steve: And, wow, thanks for the great summary and report.

Leo: What have you got there?

Steve: So I was just checking my Twitter feed on my little mini while you were talking, and I have an update about coffee. Then we'll get to BULLRUN, I promise, everyone. So this was tweeted by @ChemGuy60223. That sounds like a zip code, maybe. And so he sent two tweets. He said, "Steve, for an alternative to your coffee temp controller, check out Coffee Joulies, available on Amazon now."

Leo: They don't work.

Steve: He says, "I'm...." Huh?

Leo: They don't work.

Steve: Oh.

Leo: Well, I've used them for whiskey.

Steve: Okay, I don't think that's quite the right technology, Leo. He says, "I'm a chemist." And of course his Twitter handle is ChemGuy. He says, "I'm a chemist and have these. Use as a class demo. They really work. Uses a phase-change material to..."

Leo: It's the same idea.

Steve: "...absorb/release heat at the proper temperature.

Leo: So they're little metal beans that go in your coffee. People do these for whiskey, as well, to keep whiskey cool without watering it down.

Steve: Oh, okay.

Leo: Get the idea?

Steve: Yeah.

Leo: But the people I know who've used these have not been crazy about this as a solution, so...

Steve: Yeah, I don't want any beans. I want - yeah.

Leo: We both bought that mug, if they put it out. If they don't, we'll try the Joulies. I've been aware of these for some time. I'm not convinced they do the job. Anybody in the chatroom has a better - has experience? There's the guy in Twitter. He's a chemist.

Steve: Yeah, no, and the theory works. And so he's probably using it in order to demonstrate...

Leo: Same idea.

Steve: ...that concept in his classroom.

Leo: Right. It's a heat and energy storage device.

Steve: Yeah.

Leo: You know, it may be - I wonder if his handle is his zip code. That would be Chicago area; right? What is it?

Steve: 60223, he said.

Leo: 60223? Yeah, Midwest somewhere. Or maybe it's a chemical something. That's what I would guess.

Steve: Could be.

Leo: My college roommate and good friend just passed away at the age of 56.

Steve: Oh, how - why? So young, Leo.

Leo: I know, so young. Pancreatitis. It was unexpected. But he was the guy in Oregon for food safety and tracked down many famous cases of food safety, and he was the guy who got almonds irradiated for the future. But the main reason I thought about this is because his license plate was, and I don't know if anybody will recognize this, 0157H7, which is the deadliest strain of E. coli. It's a common culprit in food-borne illness.

Steve: Ooh, what a happy plate.

Leo: Bill was quite the character. When we were in college he used to sell T-shirts with parasites. He had a picture of a fry cook, and it said, "How would you like your eggs?" And all the eggs were parasite eggs that were common in human parasite infections. So not a best-selling T-shirt.

Steve: Yeah, that one was one that raised eyebrows.

Leo: Bill was a great man and a real loss, Bill Keene. Yeah, very young fellow. But, boy, he saved a lot of lives. So he had a good life. But in better news, continuing - oh, yeah, there's an article about him in USA Today. I didn't see that. That's good, yeah. Go ahead.

Steve: Okay. So, BULLRUN.

Leo: What is BULLRUN?

Steve: BULLRUN is the codename. And, boy, have we been learning about codenames, creative codenames that the NSA generates like crazy in the last few months, ever since Edward Snowden dropped the first of many bombs. BULLRUN surfaced about a month ago in another release of slides. And there was a slide that referred to the BULLRUN Project. And it raised a lot of concern because it discussed the NSA's active work to decrypt the Internet's encryption. And all the headlines that flashed were "NSA Breaks Internet Encryption." And of course my take at the time, and even now, is, well, okay. As Bruce Schneier has said, trust the math. And we do trust the math.

But when a cryptographer sees that, and the cryptographer spends his life thinking about these problems, it creates a psychological dilemma. It's, well, okay. If that's true, and I, a cryptographer, pride myself in teaching cryptographer and knowing everything there is to know about cryptography, how do I square that? And so that's exactly what Matthew Green did in his latest blog on Monday, December 2nd, two days ago. His blog posting - he posts in CryptographyEngineering.com is his site, and his blog post was "How Does NSA Break SSL?"

And so he had been bugged about this briefing sheet, but has been pulled away. He had intended to spend some more time on this, but has been pulled away by his own project, which has been taking up his time. Anyone looking at his Twitter feed can see that he's talking about Zerocoin, which he's very excited about. He's come up with an alternative virtual currency technology, not Bitcoin, not Litecoin, but Zerocoin, which it's on my radar, and I just haven't gotten into it yet because it's still very much evolving. So I'm going to wait for it to settle down, and then we'll doubtless talk about it.

So Matt starts out by explaining. He said: "First, I'm well aware that NSA can install malware on your computer and pwn any cryptography you choose. That doesn't interest me at all, for the simple reason that it doesn't scale well. The NSA can do this to you, but they can't do it for an entire population. And that's what really concerns me about the recent leaks: the possibility that the NSA is breaking encryption for the purpose of mass surveillance." Which of course is an entirely different problem.

So what he does is to essentially brainstorm, first sort of reasonable things, and then, a

little bit later, admittedly, more of the tinfoil hat sort of things. So we'll follow along a little bit and talk about this. The first up that Matt looks at is the concept of just outright theft, or acquisition somehow, of the raw RSA keys. And of course this approach, just the NSA getting keys is so obvious and easy that it's somewhat difficult to imagine the NSA spending much effort or resources on hyper-sophisticated attacks because we know from reports that have been published that GCHQ in Britain and our own U.S. NSA are completely comfortable and have suborned U.S. providers overseas. And even within the U.S. they've demonstrated a willingness to obtain SSL keys using their subpoena powers and gag orders.

And we, I mean, one of the observations we've made is the insane number of certificate authorities that our browsers trust. So any of those are able to mint SSL keys on demand. And who's to say that that hasn't happened? So unfortunately the whole public key infrastructure is a fundamental weakness that we talk about often on this show because that's one of the things we talk about is fundamental weaknesses. And that's like first and foremost, which he talks about.

Now, the other thing that's a little bit of a concern is the idea of, as Matt puts it, suborning hardware encryption chips. The New York Times recently ran a story where their headline was "Documents Reveal NSA Campaign Against Encryption." And the illusions in the slides which The New York Times showed, which we also covered at the time on the podcast, noted that a significant fraction of encrypted traffic on the Internet is produced by hardware devices such as SSL terminators or accelerators. The slide that we talked about a week or two ago with Google, it showed that before anything went into their own network, on their border was SSL equipment, well likely to be hardware, in order to speed this up. SSL accelerators are often what companies use on their front end of their networks to essentially deal with the otherwise high computing cost of negotiating a public key with all of the people connecting to their site. And those use hardware chips. Those hardware chips come from somewhere.

So the question is, okay, what does "suborning a hardware encryption chip" mean? And so in brainstorming, Matt suggests, he says: "The obvious guess is that each chip encrypts and exfiltrates bits of the session key via 'random' fields such as initialization vectors and handshake nonces. Indeed," Matt writes, "this is relatively easy to implement on an opaque hardware device. The interesting question is how one ensures these backdoors can only be exploited by the NSA and not by rival intelligence agencies."

So what he's positing there is that, if the hardware performs the entire handshake, the whole protocol setup, and you want that all in hardware because that's where you get the performance boost, then things like initialization vectors, which should be one-time-use nonces - they can be publicly known, but they have to always be unique - and handshake nonces, which exist in the protocol. If some of the bits from the session key were, for example, scrambled and stuck in the nonces, an observer looking at the packets would think everything was fine unless you really analyze every detail of the handshake, looking for cross-item correlation. But if the NSA knew that such-and-such a chip were used, or just the NSA knew that chips were out there which misbehaved in ways that they had participated in engineering, then that could immediately remove, for example, half of the session key length, dropping it from 128 to 64, which is then much more feasible to brute-force.

So we don't know that that's not going on. But apparently we have evidence that the NSA has been involved in this sort of pressure against encryption chips, and Matt suggests, if so, this is how it might be done. And it's chilling. And of course he's right. It isn't at all clear how you prevent that from being, I mean, this is exactly the kind of backdoor which people in the know are worried that the NSA might have installed and be

entirely relying on the secrecy of it not being broadly known. The problem is, once it's discovered, if it hasn't been discovered by others, once it's discovered, then all of this hardware that's out there is compromised. So that's horrifying and - if true. So, yikes.

Next class of attacks against SSL could be side-channel attacks. We've talked often about various sorts of side channel attacks. Something like the operating time, the resource consumption, cache timing, or radio frequency emissions are all things that have been demonstrated to leak information. Of course all of those require, typically require proximity. I mean, something like cache timing, you really can't determine the cache timing at the other end of a connection, or any outside the machine at all. But Matt notes that anytime you virtualize TLS servers, SSL servers in a cloud setting, and if spyware is able to be running in a virtual process on the same hardware, then it's sharing hardware resources. And it's entirely feasible to imagine then that malware could be watching the shared cache of the process in order to find weaknesses and try to determine the keys that are in use.

So those are very chilling. It's one of the reasons, actually, that I was very happy with the work that Dan Bernstein had done on the elliptic curve crypto that I chose for the SQRL authentication login technology because he really verified, his algorithm specifically guarantees that no secret information is involved in any timing or memory access at all. So all of the memory accesses that are done, any timing decisions, branches, never involve anything that is supposed to be a secret. So that's one of the ways you work to protect side-channel attacks.

But fundamentally we're looking at just a problem with the implementation of cryptography on platforms that are doing other things at the same time as crypto. So the argument is that's one reason that you would use hardware is that, by encapsulating, for example, the entire SSL handshake in a single chip, then nothing else can get to it. Except, as we just saw, if the NSA gets to the chip, then we're back having a problem. So the problem is all of our technology is arguably sensitive to many different types of attacks.

And another one that we've talked about is weak random number generators, which Matt brings up again. He says weak random number generators are a problem because the crypto that we're using absolutely depends upon the quality of random number generators. We use those on the client side during the RSA handshake. We've done a podcast about how SSL handshake works [SN-195]. The first thing that happens is the client generates a random value, which it sends to the server as part of its initial client hello SSL handshake. And that absolutely has to be random. It's called the premaster secret. And if an attacker is able to predict the output of that client random number generator, then it's possible to decrypt the entire session from that point.

We also use client- and server-side randomness as part of the Diffie-Hellman handshake. Both sides, remember, generate as random a number as they can, and then they use a variant of that or something derived from that which they exchange with each other. And when they each receive what the other sent, they're able to combine that and arrive at a shared secret, which they then use from then on. But that is extremely sensitive to the randomness of the content that they're each generating. And if either one of them is not sufficiently random, that can compromise the handshake technology. And, finally, long-term generation of random numbers. We've already seen where this has been a problem, for example, the generation of RSA keys that are used in server certificates. Remember that the EFF has a project called the EFF SSL Observatory.

And what we found, we the industry found, to our shock, was that there was actually duplication of private keys among many online services. They independently arrived at

certificates that were using the same private key because they were all using, I hate to think, maybe the BSAFE random number generator from RSA that was by default using a known buggy random number generator. Thus the prime numbers that they were arriving at, thinking that they were pseudorandom, weren't as pseudorandom as they believed. The consequence was that there was a lot of collision among these certificates. And so that's certainly worrying. Thus we really do need strong random number generators. And though there's no proof of it, there is some reason to believe that the NSA was acting in order to get a bad random number generator out into the NIST specification. And even though everyone who knew anything thought nobody would ever use this, turns out that that was the default random number generator in RSA's BSAFE library, across their entire product suite.

Leo: Why not? You know? Make it easy.

[Talking simultaneously]

Leo: ...all the work, you know? Let's just...

Steve: Just unbelievable. And then lastly he suggests that there could be, even though we've all fallen in love with perfect forward secrecy, for good reason, there could be esoteric weaknesses in perfect forward secrecy systems. He notes that one of the things which SSL does in order to minimize the burden of establishing the initial handshake secret is session resumption. We've talked about that often, too. The idea is that, when the endpoints already have - when the endpoints have previously agreed on a set of public key base parameters, the client will offer, essentially, a ticket of that session to the server when the client notices it's connecting to the same server it had connected to before. If the server is caching its sessions, then it'll look up, it'll use that session ticket to find the matching ticket and say, oh, and basically short-circuit the most time-consuming process.

Now, the problem is that, while that's neat, if you're just connecting to machines, what do you do if you're connecting to a server farm? Now the problem is the client could be reestablishing a connection with a different machine in a huge server farm. So now you've got to come up with a way, essentially, of moving all of the session information and session tickets out to the front, where there's some way to sort of offload that from all the individual servers. And it turns out it just creates a huge architectural headache for people who are trying to use resumption tickets, and it can be a problem.

And then Matt also notes that the Diffie-Hellman parameters, which are increasingly being used because we're all liking Diffie-Hellman as opposed to RSA for our crypto, they must be chosen with care. That is, the so-called elliptic curve parameters must be chosen with care, and that using curves which are not safe can quickly destroy the entire system's security. So in general, Diffie-Hellman is fragile and requires proper behavior from both ends of the handshake, for reasons that we've already talked about. So those are sort of the standard, not going to raise any eyebrows sorts of problems. And it's a summary of the issues that we've talked about throughout the podcast in the past.

Under what Matt considers "tinfoil hat concerns," or maybe regard them as thought experiments, the what-if sort of stuff, he talks about just what about breaking RSA keys? And Matt notes in his blog, he says: "There's a persistent rumor in our field that NSA is cracking 1024-bit RSA keys." He says: "It's doubtful this rumor stems from any real

knowledge of NSA operations. More likely it's driven by the fact that cracking 1024-bit keys is highly feasible for an organization with NSA's resources."

Okay, now, wait a minute. "Highly feasible." To get a calibration on that, back in '03, so 10 years ago, a decade ago, encryption researchers Shamir (the S of RSA) and Tromer estimated it would cost $10 million for a purpose-built machine that could factor one 1024-bit key per year. Okay, but that was 10 years ago. An updated estimate this year, in 2013, Tromer relooked at this and estimated that the numbers would be about a million dollars for cost for a machine to factor 1024-bit keys in hardware, and it might be significantly lower, which of course is pocket change for the NSA. But remember, that's a year. So we're still talking about a really substantial amount of work to do this.

But anyway, so summing it up, Matt says, why is this considered tinfoil hat? And bringing it sort of back to some reality, he says: "Because as far as we know, nobody's ever done it. Not even once. Not even at 1024 bits. So all the entire crypto community has are rough guesses." He says, since it's never been done, guesses could be dramatically too high or dramatically too low. He says and 1024-bit RSA keys - he notes that 1024-bit RSA keys are now being rapidly phased out. For example, when I renewed my certs and went to EV certs, what, about two years ago because I'm about to come up to renewing them again, I went to 2048. In fact, that was the default that Digikey was then recommending. And of course Google famously updated all of their connection technology to 2048-bit, ahead of their end-of-year deadline, which is what they were planning.

So, and we do know, what is it, 768-bit keys have been cracked. But remember that this scales exponentially. So 1024-bit is not one third harder than 768. It is dramatically, exponentially harder. And we've already doubled the key length to 2048. So everyone believes that we are secure and are going to be.

Okay. What about RC4? And this is still in tinfoil hat mode, Matt feels, the notion of cracking RC4. So it turns out that, and we've talked about cipher suites often, there's a large suite of ciphers available. Servers have their batch; clients have their batch. And they negotiate to find the strongest, or really what happens is the client gives the server the dump of all it knows about, and from among those the server chooses, in the order it wishes, what it wants.

And so, for example, GRC, my site, had been using RC4 until a few weeks ago because it was the only one that would rate me as well as I wanted to be over on SSL Labs. I didn't see any danger in it. And Matt argues that there is really no danger. But in tinfoil hat mode, he also notes that, first of all, 50% of all HTTPS traffic is still being secured with what he refers to as "creaky old RC4." And admittedly, it's starting to show its age. Remember that the attack that we covered a few months ago was done by some of the RSA guys where the way RC4 works is it uses a scrambling algorithm where it has two arrays of 256 bytes. The key sort of prescrambles the array, and RC4 generates a pseudorandom bit stream coming from the dynamic continual scrambling of the array.

So the idea is that it gets better as it goes along. But the researchers discovered, if they really looked more carefully at exactly the way the arrays were being scrambled, there was a greater lack of randomness among the choices that the system made than they previously believed. So basically they pushed the original known weakness out further, and further worried everybody about the strength of RC4. Still not to the point that it would be a real problem, but worrisome.

And so Matt writes that we don't know of any attack that would allow the NSA to usefully crack RC4. Even given all of this, with all these weaknesses, there still isn't a way, as far as we know, to usefully crack it. And he says: "The known techniques require an attacker

to collect thousands or millions of ciphertexts that are either (a) encrypted with related keys" - and that was the WEP problem that was one of the reasons we did abandon the WEP WiFi protocol - "or (b) to contain the same plaintext."

And Matt reminds us that the best-known attack against SSL using RC4 takes this same plaintext approach. It requires the victim to establish billions of sessions. This is the BEAST attack. And it's why BEAST was really never anything to worry about. And even then it only recovers fixed plaintext elements, like cookies or passwords, occurring at the front of these billions of identical queries that browsers make. So it was at worst a theoretical attack. And it's why I kept RC4 where I had it. Now I've moved it down, and we're using perfect forward secrecy. And of course everybody else is famously moving to that very rapidly because BEAST is just no - actually, BEAST, I'm sorry, BEAST was the cipher block chaining problem. Was it CRIME, I think, that was the RC4 attack? Anyway, so fundamentally, in tinfoil hat mode, we could argue that RC4 is crackable. But practically, Matt's feeling is, eh, not so much. And certainly it's in the process of leaving.

And then just out of the blue he says, well, you know, if we're going to have our tinfoil hat on, how about some sort of other side-channel attack, something that we can't imagine or don't know about, because fundamentally our cipher technology is vulnerable to those. So if you like tinfoil, you like the idea that maybe there's some way of leaking secret information that we haven't picked up on, but that the NSA knows.

So, and then finally, of course, does the NSA have secret and completely surprising quantum computing capabilities. Well, nobody thinks so. I mean, nobody seriously in academia believes that the NSA is that far ahead of us. And they would have to be radically far ahead in order to have actual quantum computing technology to crack crypto. But if we have our tinfoil hats on, who knows? And I just note that, with all of the certificate authorities that our browsers now trust, how do we know that one of them isn't the NSA? I mean, literally, a covert, from the beginning, certificate authority that is wholly co-opted, I mean, that actually is an NSA front and has a convincing set of credentials and is doing business and all of the browser vendors believe it's whomever, and it's not. I mean, that's, if I were the NSA, that's one of the things I would do. So I hope I just didn't give away one of their secrets.

Leo: Oh, do you really. Wouldn't it be funny if it were the Hong Kong Post Office?

Steve: Oh, goodness, yeah. Well, I would imagine it's something really hiding in plain sight, actually. But Matt wraps up by saying that one bright spot is that the NSA GCHQ disclosure, the disclosure that we saw, does describe their present capabilities as, quote, "extremely fragile." So that implies that the things we are already doing have significantly changed the game for them. And I think that makes absolute sense. I think that they were noting that, yes, there were things they could do, but those were fragile, meaning that a whole bunch of things had to come together in just the right way for them to be able to decrypt. And that might have been storing up traffic and then arranging to get an expired key, that we've talked about on the podcast, that kind of thing, the idea being that they recognized, if the world switched to perfect forward secrecy, so that you could no longer determine what the session key was, use the server's private key to decrypt the session key because of using ephemeral Diffie-Hellman handshake, all of that that they were using was then obsolete. And their fragile system was hurt.

So there's no question in my mind that deep within the bowels of the NSA they are not happy about - essentially the upshot from what Snowden did has clearly mobilized the

entire security industry to speed up our game, basically, take this much more seriously than they were. And because what the NSA was doing was arguably fragile, doing things like tapping Google's unencrypted links between datacenters, well, Google's going to fix that. So that's a perfect example of something fragile which has now been taken out of their grasp.

Leo: I think this was all in response to that in the first place, that they were most - the thing they were really afraid of was the already widespread use of encryption and so forth, and that they were staring at a dark Internet. And while they had great capabilities with analog circuitry and telephones, they didn't have that same kind of capability.

Steve: Well, but we were also complacent. I mean, Google assumed that their private fiber was secure. They made that assumption. And it turns out, I mean, so we all just sort of didn't recognize what was going on. And that's the other thing, too. The Snowden leaks have been so powerful because it hasn't been one fact that got out, it's been a tsunami of revelations about just how aggressive the NSA's policies have been, which has touched every corner of the security industry and has said to people, okay, we need to change things.

Leo: Yeah. Well, now what? Now what are they going to do?

Steve: Leo, there isn't anything they can do. I mean, they're going to have to appeal to the law. They're going to have to end up using subpoena powers and national security letters in order to get the things they need. They're going to have to demand the keys from companies or from the companies those companies trust, meaning the certificate authorities. I don't see any other solution. And so that means we need to adopt technologies that are purely, that are entirely TNO. We need to Trust No One. That is one of the fundamental concepts of my SQRL login technology is there is no third party. It is fully TNO. You are trusting no one else. Unfortunately, the entire Internet infrastructure with websites means we have to trust the certificate authority. So that's its point of weakness. That's where the NSA will go.

Leo: Well, you've been right on all along with your calculations about how these things are working, so...

Steve: It's because ultimately it's technology driven. And we understand technology.

Leo: Yeah, if you understand how it works, it's apparent where the attack surface must be.

Steve: Yeah.

Leo: Steve Gibson is at GRC.com. Yeah, that's the place where you can find SpinRite, the world's finest hard drive maintenance and recovery utility. You can also

find all his freebies, his Perfect Paper Passwords, all the great things he does. And while you're there, the 16Kb version of the audio, the smallest audio made available, and full transcripts so you can read and search. It's very helpful. I was searching the other day for something you mentioned. And because all those transcripts are there and online and Google indexes them, it's very easy to find what you need through those transcripts, so thank you for doing that. We have full-quality audio and video available at our website, TWiT.tv/sn for Security Now!, and of course wherever netcasts are aggregated. We have feeds for all of those, as well. Are you going to do a Q&A next week?

Steve: I don't know.

Leo: Maybe. If you've got questions for Steve, at some point we will do a Q&A. Go to GRC.com/feedback. Do not attempt to question him in any other form. Actually, I guess you're answering people on Twitter, so...

Steve: I am. And, I mean, it's a mixed blessing. I love the social network that we've established. I like being able to answer people. I feel a little bit of an obligation to do that, so sometimes it's a little disruptive when I'm deep in brain mode. But I know people will understand if I don't get back with them immediately. And for what it's worth, I mean, please do continue submitting questions to GRC.com/feedback. I would love to do a Q&A. We're driven by the news. And so if the news overwhelms us, we've got to cover the news. But we always have our listeners' questions to fall back on.

Leo: Right on.

Steve: Yeah. And a lot of, I mean, a lot of what I talk about now during the week is coming to me in real-time through...

Leo: Through Twitter.

Steve: ...the Twitter, yes.

Leo: Well, then in that case let's mention that you are @SGgrc on Twitter.

Steve: Oh, by the way, Leo, somebody, just for some reason they checked, they did a search of posts to @Sgrc. Do not do that. It turns out people have been mistweeting me to @Sgrc, and there's a whole collection of people that I never saw, and I never replied to. So don't use that. Use SG, as in obviously Steve Gibson; GRC, Gibson Research Corporation: @SGgrc.

Leo: S double GRC. And look for the guy in the fancy derby or whatever.

Steve: Nancy, my sister, took that of me over the holidays when I was up there a couple years ago, yeah.

Leo: The scarf and the beret.

Steve: Yup.

Leo: Thank you, Steve. We'll see you next Wednesday. We do this at 11:00 a.m. Pacific, 2:00 p.m. Eastern time, 19:00 UTC, every Wednesday, live, if you want to join us live and visit through the chatroom.

Steve: And next year we'll be switching to Tuesday at 1:00 p.m. Pacific time.

Leo: 4:00 p.m. Eastern, 21:00 UTC, that starts January 7th. So be prepared. Be aware. Thank you, Steve. We'll see you next week on Security Now!.

Steve: Thanks, Leo.