SECURITY NOW!

Transcript of Episode #693

Internal Bug Discovery

Description: This week we take a look at Australia's recently passed anti-encryption legislation; details of a couple more mega breaches, including a bit of Marriott follow-up; a welcome call for legislation from Microsoft; a new twist on online advertising click fraud; the DHS's interest in deanonymizing cryptocurrencies beyond Bitcoin; the changing landscape of TOR funding; an entirely foreseeable disaster with a new Internet IoT-oriented protocol; a bit of errata; and some closing-the-loop feedback from our truly terrific listeners. Then we look at a case where a prominent company discovered one of their own bugs and acted responsibly - again - and what that suggests for everyone else.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-693.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-693-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with lots to talk about. He's going to have an update on the Australian encryption law and what it could mean to a smartphone near you. We'll talk about massive hacks, of course, and a Microsoft president suggesting there should be legislation and limits on facial recognition now. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 693, recorded Tuesday, December 11th, 2018: Internal Bug Discovery.

It's time for Security Now!. Hey, better an hour late than never. Sorry to keep you waiting. Steve Gibson, he's our security guru. But as often is the case on Tuesday, we get a little backed up. Hey, Steve. Good to see you.

Steve Gibson: Yo, Leo. Yeah, you know, there are things you can do for that little backup, but...

Leo: Are you telling me there's shots for that? Anyway, sorry to keep you waiting, but welcome.

Steve: Not a problem. Glad to be back on. Well, because we've got a lot to talk about. It's funny because I guess a couple of hours ago I looked over in Twitter, where I had not looked for a while, and I was overwhelmed by the quality of the incoming tweets. And as it was, we already had, like, too much to talk about. So I thought, okay, shoot. But I'm going to spend some time over there and catch up because we just have so many great listeners who are doing a great job of finding stuff and making comments and providing feedback.

I wanted to talk about, today, we titled this podcast number 693 for the 11th of December "Internal Bug Discovery," which was motivated by a posting made by a prominent company in the last week who discovered a rather significant privacy breach in their own API and dealt with it. And what this put me in mind of is that there are four different sources of discovery of bugs, four different sort of like discoverers. And only one of them is really the way you want the world to work, which is what we just saw. But it's also the most expensive of the four for a company. So anyway, I just wanted to spend some time and sort of get a little philosophical about this aspect of the world which is becoming more important, after we catch up with a lot of security.

We've got, of course, the most tweeted topic was this Australia recently passed anti-encryption legislation. We also have details of a couple more mega breaches, including a bit of Marriott follow-up. A welcome call for legislation from the president of Microsoft, of all places. A new twist on online advertising click fraud that just made me chuckle. The DHS, the U.S. Department of Homeland Security, has put out a document asking for proposals for the possibility of developing a range of technologies, one of which is doing, I was going to say four, but maybe it ought to be two, the other cryptocurrencies, what has been done to Bitcoin in terms of deanonymizing it in order to track people. We want to touch on that.

We've also got the welcome and changing landscape of Tor's funding, The Onion Router anonymizing network funding. An entirely foreseeable disaster with a new and forthcoming Internet IoT-oriented protocol. We have a bit of errata, some closing-the-loop feedback from our truly terrific listeners, and then we'll take a look, as I talked about at the beginning, a look at this case where a prominent company discovered one of their own bugs and acted responsibly and what that means. And we have just a very fun Picture of the Week. One of our listeners spotted this, took a picture, sent it to me with the caption that I'm using, essentially. So I think we'll have fun sharing that, as well.

Leo: That's funny, yeah.

Steve: So another great podcast, I think.

Leo: Awesome.

Steve: So anyway, our Picture of the Week is just a kick. It's a photo that one of our listeners took of the back door of LogMeIn's location.

Leo: Oh, that's funny.

Steve: We know that because there's the LogMeIn logo prominently shown, and then the door below that says "USE MAIN ENTRANCE" with an arrow pointing to the left. And so this was sent to me with the caption, "Apparently LogMeIn has a backdoor."

Leo: Okay.

Steve: Okay.

Leo: They do not have a real backdoor. Well, they do have a real backdoor. They don't have a virtual backdoor.

Steve: Just to be clear, yes. Their backdoor is in the physical world, not in the cyber world.

Leo: Yes.

Steve: So unfortunately, speaking of backdoors, Australia's Telecommunications Assistance and Access Bill...

Leo: Oh, yeah, I was wondering what you'd say about this.

Steve: Yeah, of 2018. Last Thursday Australia's House of Representatives has finally passed what is known as, formally, the Telecommunications Assistance and Access Bill of 2018, less formally known as the anti-encryption bill. And once Australia's upper house votes the bill into law, which is expected, since the bill enjoyed wide bipartisan support, it would go into effect during the next session of Australia's Parliament, which occurs in early 2019. So this looks like it's happening.

It would provide, and I'll go into some, you know, this podcast is about interesting details, so I want to cover a little bit of probably what the regular mainstream press doesn't talk about. It would, as we would expect, I mean, and we've been talking about this sort of thing happening now for quite a while, provide law enforcement with a legal means, that is, lawful means to compel anyone using encryption - so Google, Facebook, WhatsApp, Signal, and others - to assist them in accessing the encrypted communications of their products and platforms.

So the way this thing is written, they use the term "notice" in an unusual way, but it takes the form of three so-called "notices." There's the TAR, the T-A-R, which is the Technical Assistance Request. And that's described as a notice to request tech companies for providing voluntary assistance to law enforcement, which includes removing electronic protection, providing technical information, installing software, putting information in a particular format, and facilitating access to devices or services. Okay, so "voluntary" is like the key word there. That's the Technical Assistance Request.

Then there's the TAN, the Technical Assistance Notice, which the bill states is a notice "requiring" rather than "requesting" technology companies to provide assistance they're capable of providing that is reasonable, proportionate, practical, and technically feasible, giving Australian agencies the flexibility to seek decryption of encrypted communications in circumstances where companies have existing - and that's key - existing means to do it. Like, for example, in the bill it states "points where messages are not end-to-end encrypted," meaning either before they have been or after they have been decrypted at the other end. And then the TCN is the so-called Technical Capability Notice, which is issued by the Attorney-General, requiring companies to "build a new capability" to decrypt communications for Australian law enforcement. So we can all read that.

Collectively, these so-called "notices" would compel tech companies to modify their software, where modification is necessary, and their service infrastructure to essentially backdoor encrypted communications and data that would otherwise not be obtained. The bill itself I looked at, and it is incomprehensible to a lay technologist. I mean, it is just - I think it was 218-some pages. However, I found a sort of a digest of it which is more conversational and has some back-and-forth and breaks these things down.

So I want to share a couple things because, for example, Apple was involved in providing feedback, and the industry was invited to provide that feedback, although all of this has been relatively fast-tracked. I mean, it's as if suddenly there's a great deal of urgency. And a lot of the opposition to this, and you can imagine there's plenty from the civil rights people and the industry itself that is saying, whoa, hold on, let's not rush into anything.

And the other problem we have is that there is, as always seems to be where legislation bumps into technology, there is ambiguity left in the legislation which of course ultimately ends up needing resolution in courts when someone says, "Well, we don't think that's what the bill says," and the people who wrote it say, "Well, yes, but that's what we meant." And then the people opposing it say, "Well, then, why didn't you say so?" Well, the reason they didn't say so is they know they couldn't have gotten it passed if it had been a lot more explicit. So it's deliberately left in the gray in order to get it through the legislative process and dump it at the feet of the courts.

So from this really good digest, they said: "Following earlier industry consultations, the Government released an Exposure Draft" - which is what they called it - "of the Bill on 14 August 2018." Okay? Right? So August, September, October, November, December, not a long time ago - "and sought public submissions by 10 September 2018." So whoops, wait a minute. August, September. So public submissions within shy of four weeks.

"The Department of Home Affairs [in Australia], DoHA, received almost 16,000 submissions, of which over 15,000 were classified as standard campaign responses" - so I guess politically motivated - "743 were [described as] unique individual responses classified as appropriate for consideration, and 55 were considered substantive submissions from industry groups, civil society, government bodies and individuals."

They wrote: "While some stakeholders raised concerns about other schedules" - oh. This thing has, I think it was four, they called them "schedules," which are essentially the meat of the legislation. It's Schedule 1 which is where all of these three notices were described. So this says: "While some stakeholders raised concerns about other schedules, the majority of submissions focused primarily or exclusively on Schedule 1 of the Exposure Draft." And that is the so-called "industry assistance" schedule.

They wrote: "Many stakeholders provided submissions that included general and specific recommendations on the proposed industry assistance scheme, including" - and I don't know what these acronyms are - "IGIS, AHRC, LCA and applied cryptography academics" - so those must be Australian interest groups of some sort - "and cryptography academics Chris Culnane and Vanessa Teague. There was significant concern that the scheme in its current form has very wide application, and that amendments to offer greater definition, narrow the scope, or clarify processes are necessary.

"From a technology perspective, Apple submitted that Schedule 1 'remains dangerously ambiguous with respect to encryption and security.' Further, Apple stated: 'We encourage the government to stand by their stated intention not to weaken encryption or compel providers to build systemic weaknesses into their products.'" This is still Apple talking.

"'Due to the breadth and vagueness of the Bill's authorities ... the Bill could allow the government to order the makers of smart home speakers to install persistent eavesdropping capabilities into a person's home, require a provider to monitor health data of its customers for indications of drug use, or require the development of tools that can unlock a particular user's device, regardless of whether such tool could be used to unlock every other user's device as well.... While we share the goal of protecting the public and communities, we believe more work needs to be done on the Bill to iron out the ambiguities on encryption and security to ensure that Australians are protected to the greatest extent possible in the digital world.'"

And none of that happened. So this thing was pushed through and rushed, and the bill as it stood is what is in the process of being enacted. So for its part the government argues the new legislation, naturally, we've heard this before in various forms from multiple governments, "for national security and an essential tool to help law enforcement and security agencies fight serious offenses such as crime, terrorism, drug trafficking, smuggling, and [of course it's in there] the sexual exploitation of children."

Since the bill had support from both major parties, the Coalition and the Labor parties, the upper house could vote in support of this Assistance and Access Bill to make it law, which as I said at the top is expected to come into effect in early 2019. The bill states that the tech companies cannot be compelled to induce a systemic weakness or systemic backdoor into their legitimate software or hardware, or to remove electronic protection like encryption to satisfy the government demands. So again, here we have this weird collision of nontechnical intent and the statement of what the bill cannot be compelled to cause companies to do, yet they're being compelled to essentially break encryption in some form or fashion. So, okay. So who knows how this is going to end up coming out.

"Instead, the new legislation contains measures aimed at facilitating" - this digest says - "lawful access to information through two avenues, decryption of encrypted technologies [okay] and access to communications and data at points where they are not encrypted." So I guess somebody has told these legislators that, oh, it's not necessary to decrypt the data if you just get it before it's encrypted or after it's decrypted.

So the bill stipulates, quoting from the bill now: "We encourage the government to stand by their stated intention" - oh, I'm sorry, that's what Apple had said before. So we have, as we've said before, this Five Eyes Nations and the so-called "going dark" problem that we've talked about. And I've looked around at what various technical outlets were having to say about this. One said that since Australia is a member of the Five Eyes alliance, along with the United States, the U.K., Canada, and New Zealand, which last month declared that "privacy is not an absolute" and the use of end-to-end encryption "should be rare," the new bill could be a stepping stone towards new encryption laws in other nations, as well.

I mean, and that's what we're seeing. We're seeing everything, well, look at the GDPR. We're seeing everything sort of creeping forward. Australian Prime Minister Malcolm Turnbull has previously made his position on encryption clear, last year saying that the laws of mathematics are very commendable, but the only law that applies in Australia is the law of Australia. And Apple at some point responded to the new bill by making their submission to the Australian government, saying that encryption is simply math, and any process that weakens the mathematical models that protect user data for anyone will by extension weaken the protections for everyone.

So anyway, I really do think that this is the interesting thing, I mean, this will be probably the most interesting issue in 2019, given that Intel's processors have now been fully examined and don't have any more disasters waiting for us, and surprises, because even as we've talked about, in the U.S. there is this notion of warrantable examination. In the U.S., privacy is not an absolute. You need to get a search warrant, in which cases you're able to search under that warrant. And we do have a philosophical collision here between math, which when properly implemented, as Apple says, is absolute. It can provide absolute encryption that a third party, you know, it's unlike the front door of a house that can be broken into.

So anyway, it just looks like this is a battle between governments and the forces that want to protect our privacy. And I expect governments to win because they're able to ask for what they want. And Leo, I like your take when you talk about that the amount of visibility that now exists as a consequence of the incredible use of electronic communications compared to decades ago. And your analogy, I think, is really good about a few dark pixels on a high-res screen.

Leo: And they don't like it.

Steve: And they don't want any pixels that prevent them from seeing absolutely everything they want.

Leo: That's Phil Zimmermann's - the creator of PGP - analogy. And he actually coined Zimmermann's Law that, with the advance of technology, surveillance would advance comparably.

Steve: Ah.

Leo: We were talking about the "Bodyguard," that four-part or six-part episode on...

Steve: Oh, yeah, yeah, yeah.

Leo: From the BBC. And one thing I noticed watching the "Bodyguard," and I think it's accurate, especially in the U.K., the amazing amount of information they can get just because everybody's on camera.

Steve: Yes.

Leo: Pull the camera for that. Pull the camera from that. And I think that that's a lot of what police work is these days. Well, let's see the video from last night because we have it. We've got eyes everywhere.

Steve: Right. Right. And, I mean, it's now become a staple of TV and movies where the detective goes out to the scene and then slowly does a 360, looking for all the cameras that happen to have that region in their field of view, and then says to their underling, okay, go pull the tapes.

Leo: Go get those, yup.

Steve: Yeah, exact.

Leo: So, I mean, talk about surveillance. And then you've got all of the information that you can get from a phone that the carrier has, including location information. The encrypted messages seems to me to be the smallest part of the information.

Steve: Well, and in fact we've talked about this in other contexts where we might call everything else the "metadata." Like the envelope on the letter. And sure, you may not know what's in the letter. But the fact of the metadata tells you who's talking to who. You get to build a whole network of intercommunications. And as we know, none of that is protected. It's only the content. And sure they'd like to know what's there. But even absent that, it's a surveillance dream for law enforcement.

Leo: And as you and I have both said, I mean, of course we want bad guys to be caught. We don't want terrorists to act with impunity. But they have a lot of information, A. And the bigger issue is that any compromise of the encryption integrity of any of these devices is going to not only let law enforcement read it, but bad guys will get access to it, too. And who knows what the security of a phone would be like if it's been compromised like that. It's just...

Steve: Yeah.

Leo: It would be bad for everybody. I don't know what the answer is. You at one point thought that it would be possible, your position on this was it would be possible to somehow let government to get access to this without compromising everybody's safety.

Steve: So in a very selective ecosystem, which is Apple's. And I use them as an example because it's in place, and it exists. It would be absolutely possible for Apple to curate a master key for each of their devices. Which does create a concentrated massive region of responsibility. But we know how to protect secrets. And, I mean, that's what Apple will end up doing. If this legislation happens, then the model on the iOS platform will change. And when the device is created, I mean, everything's already in place. Devices have unique secrets. They've got unique keys. They've got Secure Enclaves. All of the infrastructure is in place. All that has to happen is that, under encryption, the device sends an additional unlock key to Apple, who then stores each of its customers' devices unique unlock key. It doesn't weaken anything. All of the technology is there except it does create, inarguably, a single point of failure and a great deal of responsibility.

But given that, what it means is that Apple could respond to a subpoena for a specific device so that, for example, after the terror attack that generated so much news here in Southern California for Farood, was it? Farook? I don't remember his name now. You know, the guy who had the cell phone that could not be cracked. Well, the FBI could generate a subpoena, and they would say to Apple, we need this one phone unlocked. Apple looks up the unique key for that one device and unlocks it for law enforcement. In no way, except that Apple maintains a master key for each of their devices...

Leo: So they'd have to keep that secure, but I think they could do that; right?

Steve: Yes. And that's my argument is that, you know, I get the theoretical ivory tower, there's no way to do this without installing a backdoor or weakening security. That's just not the case. I mean, with the under...

Leo: Here's the problem. You could say, oh, yeah, well, we could trust Apple. But then there's Huawei and Samsung and...

Steve: That's why I started, yes, that's why I start off by saying this only applies to Apple's ecosystem because this is the way Apple has built this out. You're right, it's entirely - like for example the Android platform would be, you know, that doesn't exist because you've got all...

Leo: Each manufacturer would be a separate issue; right?

Steve: Yes. And they're not even updating these devices.

Leo: No, no.

Steve: I mean, that's like, you know, it's nuts.

Leo: So this Australian law is nothing like that, really. It's requiring everybody to have some way to get that data.

Steve: Yes. Yes. And, I mean, from a technology standpoint we will be talking about this, I think, because what will happen? Now, in some of the dialogue that I've seen that I didn't discuss, there has been the discussion of major - I think there are also very steep fines. I think it was $10 million is one of the fines that I saw referred to. Yet it was noted that, well, that doesn't put a dent in one of these massive multinational technology companies, so it's like, eh. And they could react by saying, fine, we'll pull out of Australia. Well, that hurts Australia, not to have the services of a company that says we either protect the privacy rights of our customers who may happen to be Australians, or we're not going to have any Australian customers.

So, I mean, but on the other hand, what have we seen before? We've seen Google initially pull out of China when China said, you know, we need you to filter and play by our rules. Google initially said no thanks and pulled out. Well, then they came creeping back in a few years later. So I think, ultimately, I do believe technology can provide answers. And I get it that the ivory tower cryptographers would rather say it's impossible. But nothing is impossible for technology, for math. So I imagine we'll end up finding some solutions. And it's going to be a great topic for Security Now!.

Leo: Well, we'll see. As you say, we'll be reporting on this more in future.

Steve: Yeah. So one of the questions is how do we get companies to be more responsible? And I'd noted over the past week that the attorneys general of 12 states - Arizona, Arkansas, Florida, Indiana, Iowa, Kansas, Kentucky, Louisiana, Minnesota, Nebraska, North Carolina, and Wisconsin - have gathered together to file a lawsuit against an Indiana-based medical informatics engineering company. Well, in fact their name is Medical Informatics Engineering, MIE. And they have a subsidiary, NoMoreClipboard, NMC. And they're a company that is automating medical records management, as NoMoreClipboard sounds like, so that instead of having all paper records, everything's online.

MIE, the parent company, sells this web-based electronic health record services. And no one has to be told that you have to take security seriously if you're going to do electronic medical records in the cloud because the cloud is the cloud. It's very different from a local network in a doctor's office, where everything is kept local. So the wheels of justice turn slowly. This all began back a little over three years ago, on May 7th of 2015, when hackers stole very personal information of 3.9 million people from MIE's backend database using a simple SQL injection attack.

And of course being medical records, it's not just LinkedIn information. It's names, addresses, Social Security numbers, also health data including lab test results, health insurance policy information, past diagnoses, disability codes, doctors' names, medical conditions, the names and birth certificates and birth statistics of their children - in other words, everything. So this joint complaint accuses MIE of failing to properly secure its computer systems, not telling people about its system weaknesses of which it was aware, and then failing to provide timely notifications of the incident.

MIE never bothered to actually encrypt this information, even though it said it did, so it was in violation of HIPAA regulations just there. It also provided public accounts for sharing the passwords "tester" and "testing," which were established specifically so that a client's employees did not have to log in with a unique userID. Yeah. Why bother with that? Penetration testers uncovered the issue and highlighted the risk to MIE, but the lawsuit says that MIE took no action. And then, using one of these test accounts which was discovered, the thieves explored the health record database using SQL injection attacks to gain further access to more privileged accounts, so using the technique we've described in the past, pivoting from one access point to another.

MIE allegedly didn't have any data exfiltration alarms in place. Nothing alerted them to the fact that data was being taken out of their system. It was a network performance monitoring alarm, well, of sorts, that finally raised the red flag because the attackers were dumping so many large records from the database at such volume that it choked off the network bandwidth, and they started to wonder, hmm, why is the Internet so slow today? Uh-huh.

Anyway, the states which have brought this suit allege that once the breach was discovered, MIE only had a draft incidence response plan, and that there was no evidence that it even followed that in any case. And they added that the notifications, had they been followed, were inadequate. They discovered the breach on the 26th of May in 2015 and informed the public of the breach via a notice on its website, not directly, 15 days later, on June 10th; and then finally began email notifications another five weeks after that on July 7th; and then finally followed up with printed paper letters in December.

So anyway, the 12 states that are bringing the suit claim that MIE and their subsidiary, NMC, violated federal HIPAA legislation protecting the privacy of health information; also accused MIE of breaking 27 state-level laws concerning data breach notification, abusive and deceptive practices, and personal information protection. The states are proposing a consent decree to clear up the matter before getting into litigation. And they're calling for an as-yet undefined payout from MIE, along with its commitment to clean up its act and follow several security measures, including the use of multifactor authentication, not making generic accounts accessible via the Internet, using strong passwords, training their staff properly in cybersecurity, using a security incident and event monitoring solution, and putting SQL injection attack detection measures in place. And they'll have to do responsible things like conduct regular security audits with help from a qualified professional, file reports, and take actions on them. In short, the settlement asks the company to do what any competent cybersecurity team charged with protecting sensitive data should be doing.

So in their coverage of this, Sophos notes the collaborative nature of the settlement, and they said: "As voices call for stricter federal privacy protection laws, this could be a sign that states are getting fed up with these mega breaches and are taking things into their own hands." They noted that in October Uber settled with all 50 states over the handling of its 2016 data breach, paying $148 million. So Sophos wonders whether this suit might herald more coordination between attorneys general to hold companies accountable.

And for our part, let's hope that other companies observe this. We don't know yet what the fine is going to be. But, I mean, this is a company clearly in violation of existing laws who has been caught doing this. And I think the fine needs to be hefty so that other companies look at that and think, whoops, it would be much less expensive for us to follow the law and put these measures in place than it would be to be caught not doing so and then have to pay the fine. Yikes.

And speaking of databases, the director of Cyber Risk Research at Hacken, Bob Diachenko, has been tracking a publicly accessible instance of MongoDB, which could be accessed without authentication, for some time. Over the course of several months, October and November, Bob initially discovered this database which was open in October, containing 66,147,856 unique records containing the full name, personal or professional email address, user's location, details of their skills, a phone number, employment history, also a link to the individual's LinkedIn profile was present, all of which led him to believe that this was likely scraped data from online LinkedIn profiles. He was unable to determine the owner of this database, just having found it at some IP on the Internet and unable to attribute it to whomever. But it is now no longer up.

He had an interview with Bleeping Computer where he noted that the scraping of personal data, presuming that it was just scraped from the web, is legal; and that making a copy of it publicly available is legal, as long as it's not used against the best interests of its owner, which is considered an offense. There were, as I mentioned, 66-plus million records, including email addresses. That entire database has been uploaded to Troy Hunt's HaveIBeenPwned service. So any of our listeners who are interested could make another visit over to Troy's HaveIBeenPwned service, put any of their email addresses in to see if they are newly apparent, which would indicate that this data had gotten loose since the last time they checked.

And, you know, I think what we're seeing is that this is a - things like this, incidents like this, because it's not the first time it's happened, it's not going to be the last time, it's a natural consequence of there now being a global network of commercial entities that are wanting to share and exchange data. And that, coupled with the incredible dropping cost of long-term high-volume mass storage, which means that everybody can have a copy of everything, and the inherent resale value of aggregated personal information on millions of individuals. And as has been observed before, we're no longer the consumer who is buying, but rather the details of our existence is the product which is being sold, completely without our knowledge or permission. You know, we've talked about Equifax that has all of this data that they have vacuumed up and assembled and then are reselling for their own profit, completely without our knowledge or permission.

So I get it that many people will not and do not care about the circulation of their personal information. But for those who do, I think the only recourse we have is minimizing what is put online where possible, and in many cases it's not possible to do that. But also, for example, in the case of the firms that provide us the ability to lock access, like the large credit clearinghouses, preventing third parties from using them to acquire credit in our name, thus essentially effecting a form of identity theft, really does make possible. You know, Leo, you and I have talked about this, the idea of locking access if we ourselves are not applying for credit actively. I'll say again, it really makes sense.

Leo: The right thing to do, you bet.

Steve: To keep anybody from applying for credit. So this was interesting. On the 6th of this month, Brad Smith, the president of Microsoft, posted another piece sort of on something he's been thinking about, clearly. This one was titled "Facial Recognition: It's Time for Action." But what was interesting was that this, as I said, was the second of two. And we didn't talk about this. When he brought it up the first time back in the summer, July 13th, was sort of his foundation-laying statement. And that's the one I wanted to focus on because he's just sort of reiterating that position last week. But on July 13th he said, and this is what I think is interesting, he says: "Some emerging uses are both positive and potentially even profound. Imagine finding a young missing child by recognizing her as she is being walked down the street." And Leo, of course, this connects right into our discussion of cameras for surveillance a second ago.

Leo: They're even more valuable if you've got face recognition on them; right?

Steve: Yeah. And he says: "Imagine helping the police to identify a terrorist bent on destruction as he walks into the arena where you're attending a sporting event. Imagine a smartphone camera and app that tells a person who is blind the name of the individual who has just walked into a room to join a meeting.

"But other potential applications are more sobering. Imagine a government tracking everywhere you walked over the past month without your permission or knowledge. Imagine a database of everyone who attended a political rally that constitutes the very essence of free speech. Imagine the stores of a shopping mall using facial recognition to share information with each other about each shelf that you browse and product you buy, without asking you first." He says: "This has long been the stuff of science fiction and popular movies like 'Minority Report,' 'Enemy of the State,' and even '1984,' but now it's on the verge of becoming possible."

He says: "Perhaps as much as any advance, facial recognition raises a critical question: What role do we want this type of technology to play in everyday society?" And he has a separate topic then, the need for government regulation. He says: "The only effective way to manage the use of technology by a government is for the government proactively to manage this use itself. And if there are concerns about how a technology will be deployed more broadly across society, the only way to regulate this broad use is for the government to do so. This in fact is what we believe is needed today" - this is Microsoft speaking - "a government initiative to regulate the proper use of facial recognition technology, informed first by a bipartisan and expert commission."

And he says: "So what issues should be addressed through government regulation? That's one of the most important initial questions to address. As a starting point, we [Microsoft] believe governments should consider the following issues, among others." And these are the bullet points that I wanted to bring up.

He said: "Should law enforcement use of facial recognition be subject to human oversight and controls? Should restrictions on the use of unaided facial recognition technology as evidence of an individual's guilt or innocence of a crime? Similarly, should we ensure there is civilian oversight and accountability for the use of facial recognition as part of governmental national security technology practices?

"What types of legal measures can prevent use of facial recognition for racial profiling and other violations of rights while still permitting the beneficial uses of the technology?

Should use of facial recognition by public authorities or others be subject to minimum performance levels on accuracy?

"Should the law require that retailers post visible notice of their use of facial recognition technology in public spaces? Should the law require that companies obtain prior consent before collecting individuals' images for facial recognition? If so, in what situations and places should this apply? And what is the appropriate way to ask for and obtain such consent?

"Should we ensure that individuals have the right to know what photos have been collected and stored that have been identified with their names and faces? Should we create processes that afford legal rights to individuals who believe they have been misidentified by a facial recognition system?"

So from a technology standpoint, which is of course the approach we take primarily, I think these are really interesting points. And you know, for example, Leo, that we've all seen signs when we enter a retail establishment, a notice like "video surveillance in use" sort of thing, which is taken to a different level if, it seems to me, if it's facial recognition, automated facial recognition running behind those surveillance videos. I mean, if you look around in any department store or in any large public facility, you often see those little black domes that are presumably cameras pointing somewhere. And it's different to imagine that they're going into, being spooled on hard drives in case there's some reason to later run back at some point and see what happened somewhere, in order to recreate an incident that happened in a public place. That seems different than them running facial recognition and logging the names, well, more than the names, the unique identities of the people that the camera believes it has seen within its field of view on a streaming basis.

But as Brad notes, this really is where we are. I mean, this is not science fiction any longer. It's probably already happening somewhere, and we don't know about it. So we were once upon a time, like I guess it was just after Edward Snowden's revelations, it came as some shock that the NSA had mega nodes on the Internet where they were sucking up all of the unencrypted packet traffic, well, even the encrypted traffic, assuming they could decrypt it in the future. And that was a little bracing for us. And so now we're in a place where, thanks to computational capability, the crazy drop in prices of sensors and processing which makes this feasible, that there could be recognition happening pervasively in the public sphere, not just on the Internet. Frightening.

So I got a kick out of a bit of news of some ad click fraud which was going on in a set of Android apps. A long time ago, somewhere here or on the TWiT network, we noted the interesting fact, I remember this being discussed, either I was discussing it with you, Leo, or it was being discussed on one of the podcasts, the interesting fact that in this very tightly optimized online advertising LAN, ads clicked by iOS device users were considered to be more valuable than those clicked by Android users.

Leo: Yeah, absolutely, because they spend more money.

Steve: Exactly. And so consequently, advertisers were shelling out fractionally more money for iOS clicks than clicks from non-iOS devices. So it should come as little surprise that, since there's already no honor among thieves, Sophos recently uncovered 25 apps on the Google Play Store. They're in the Google Play Store, so they're infecting Android platform, but they are spoofing what is already their highly spoofable user-agent headers to lie about the device.

Leo: Oh, golly. Of course.

Steve: Of course. Of course. Whose user was supposedly, though not actually, clicking on web page ads, and of course they are claiming to be various iOS devices rather than the devices they actually are because it's already ad fraud, so why not get paid more for the spoofed click on a ad? You know? Of course, why not? So anyway, I just got a kick out of that. The users, the Android devices are already blissfully unaware because this is happening offscreen that these clicks are being sent.

Leo: This is a really valuable click, honest.

Steve: Yeah, exactly. So it's from somebody on an iOS platform, so pay us more. Wow. Anyway, that just was a quickie, but I got a kick out of it. And we talked about the big Marriott breach as a consequence of their Starwood property acquisition. It just came out that a Marriott spokesman said the hotel was working on a way to reimburse some of their prior guests the cost of obtaining replacement passports.

Leo: Oh, my.

Steve: Yeah, if they can show that they've been victims of fraudulent operations where the passport number was involved. It turns out that New York Senator Chuck Schumer publicly called on Marriott over the weekend to pay for people's passport replacement fees. He said: "The data breach at Marriott compromised millions of travelers' U.S. passport info, and a new passport costs $110. So Marriott must personally notify customers at greatest risk, and Marriott should pay the costs of a new passport for victims who request it."

Whereupon Connie Kim, a Marriott spokesperson, wrote in an email to the Washington Post: "We are setting up a process to work with our guests who believe that they have experienced fraud as a result of their passports being involved in this incident. If through that process we determine that fraud has taken place, then the company will reimburse guests for the costs associated with getting a new passport." So there. Anyway, yeah. It was a big breach. And as we know, it happened four years ago. Marriott acquired Starwood, this collection of Starwood properties about two years ago. So the breach happened not directly under their control, but maybe they didn't fully vet the properties that they acquired. Who knows what the back story is. But it's nice that they're stepping up and being responsible.

Meanwhile, the U.S. Department of Homeland Security hopes to de-privacy-ify, if that's a word, the explicitly private Zcash and Monero currencies. There's this thing known as a pre-solicitation document which - and I have a link to it here in the show notes for anyone who's curious for all the details. And it also wants some other interesting technologies, like AI for image recognition of potential explosives in luggage. So the idea would be you would run luggage through some imaging scanner, and they're wanting to see if they can develop AI to potentially - and I would imagine a human would get brought in, depending upon how certain the AI was that an image might represent something explosive.

But there's a whole bunch of cool stuff in this pre-solicitation document. I only wanted to talk about the last one on the list of items. They're looking for the same sort of transaction information for Zcash and Monero that they have been able to obtain for Bitcoin. As we know, law enforcement has had great success in tracking Bitcoin transactions because, although the Bitcoin system cryptographically protects itself and is secure against external tampering, you know, Bitcoin transactions are not private since they are recorded in a public ledger, but it is tamperproof.

So the whole point of it is there's no way to go back and change the public ledger, but the transactions themselves show publicly available Bitcoin addresses, so everyone can see where the money is going to the level of the Bitcoin address. Not so with the deliberately privacy enhancing Zcash and Monero currencies.

So this DHS pre-solicitation document begins by explaining: "The Department of Homeland Security (DHS) Small Business Innovation Research (SBIR) Program, comprised of the Science and Technology (S&T)" - we're big on acronyms here - "Directorate's SBIR Program and the Countering Weapons of Mass Destruction (CWMD) Office SBIR Program" - and remember that's the Small Business Innovation Research - "invites small business concerns to review this pre-solicitation notice, which is intended to lead to the FY19" - that would be Fiscal Year 2019 - "DHS SBIR Phase I solicitation. This notice is not itself a solicitation or Request for Proposals. This notice is merely an opportunity for interested parties to comment on, or request information about, the attached topic areas."

And as I said, there are a list of topics. The last one is what caught my attention. The last item on the list says "Objective: Design a product to support the implementation of blockchain-based forensics, data analysis, and information sharing." And under description they said: "Blockchain and Distributed Ledger Technology (DLT) are emerging technologies being leveraged for a wide range of commercial and governmental applications. The most well-known use case would likely be Bitcoin within the newly emerged cryptocurrency arena, which has spurred further interest and developments. Prior efforts have addressed Bitcoin analytics, which covers only a limited scope within the greater realm of cryptocurrencies. This proposal seeks applications of blockchain forensic analytics for newer cryptocurrencies such as Zcash and Monero. And ongoing research within the field also contributes to new technological implementations and technologies that continue to multiply the specific types of consensus, privacy, security, and proof mechanisms."

They said: "A key feature underlying these newer blockchain platforms that is frequently emphasized is the capability for anonymity and privacy protection. While these features are desirable, there is similarly a compelling interest in tracing and understanding transactions and actions on the blockchain of an illegal nature. To that end, this proposal calls for solutions that enable law enforcement investigations to perform forensic analysis on blockchain transactions. This analysis can be approached in any number of ways and may consider different data situation use cases depending on whether additional data from off-chain sources are available. Furthermore, with the proliferation of new blockchain variants, the desired solution should either attempt to show generality or extensibility, or at least provide working approaches to treating newer blockchain implementations."

And I had in the show notes here, but I won't go through it, different details. They break this into three phases: the design of an analysis ecosystem where they mention Zcash and Monero; phase two, the prototype and demonstration of these forensic technologies designed during phase one; and then phase three for commercial or government applications. So it doesn't take much reading between the lines to see that law enforcement is a little unhappy that the newer non-Bitcoin so-called distributed ledger technologies are thwarting their ability to see into the transactions to the same degree that they're able to see into Bitcoin transactions.

In the case of Bitcoin and the tracking of those transactions, thanks to the help of the private sector, U.S. law enforcement authorities, for example, were able to determine that 95% of all ransomware ransom payments were cashed out and converted into fiat currency through the BTC-e cryptocurrency exchange. n international arrest warrant for the owner of this BTC-e portal was issued. But the U.S. lost the extradition fight with Russia. So the point is that, in the case of Bitcoin, it is the case that we have the technology, and the government is using forensics tools to, as we know, track where the money is flowing with bitcoin. They very much want to be able to do the same thing with the next generation of emerging cryptocurrencies such as Zcash and Monero. But at this point they don't have those tools. So they would like to have them.

And just I thought that was interesting, that this is now, not surprisingly, on their radar, and an opportunity for some enterprising individuals to see if they can help the government to produce those tools, or produce those tools for the government. It's always been a little bit of a concern to especially people who are skeptical, that the Tor Project was created by DARPA, the Defense Advanced Research Project Agency, and that its initial funding and its funding for quite a while through its startup was from the U.S. State Department. It's like, okay. But on the other hand, its technology has been vetted. It's open. The source is open. It's been scrutinized extremely well.

And so I really do think - we know that the people in the Tor Project, we know that their hearts are in the right places. And just to step back for any new listeners, what we know is that the Internet, in the same way that it was never designed with secrecy and privacy in mind, thus the underlying protocols do not support encryption, that had to be added on top later. Similarly, the Internet was never designed with anonymity in mind. We've talked about it often how the system is inherently a point-to-point system. Your browser needs to contact Microsoft, so it asks your DNS server for the IP address of Microsoft.com. And when it gets it, it sends a packet to that IP address. The source of the packet is your IP so that answers can come back to you. And the destination is the IP that it got from DNS.

And so what's established over this amazingly heterogeneous network of linked routers is essentially a point-to-point communications link where each end knows the address of the other. In other words, zero anonymity. And for what it's worth, anybody looking at packets can also see their source and destination addresses, even when they're encrypted. When the packet itself contains an envelope whose contents is encrypted and cannot be seen, the addressing information, as we were talking about earlier, the metadata is the source and destination IPs which are out there for everyone to see because the routers have to have it in order to send the data back and forth between the two parties.

So Tor came along as a unique and truly cool question, which is, in the same way that security and privacy has been added on top of the existing Internet, can anonymity be somehow added on top of an existing routing system that doesn't itself support anonymity? And if anyone is interested, we did a really cool podcast way back in the day on Tor. It was called "The Onion Router," which is what the initials T-O-R originally stood for, about the technology that solved this problem to a very, very good degree.

Tor added an anonymity layer on top of the Internet such that somebody who very much wanted to communicate with some other entity on the Internet, but very much wanted to protect their identity, was able to choose a bunch of routers that their data would bounce around, Tor routers, before emerging onto the public Internet, and would create this so-called "onion," a series of encapsulations, successive encapsulations of their data such that each router that their data went to could only take the outer wrapper off before forwarding it. The router would know where the data came from and where the data was going to, but not where it came from before that, and not where it was going to after that as a consequence of the successive encapsulations. And so after a few hops the actual source of the data is deeply obscured.

And subsequently we've talked about various academic attacks on this Tor network and how that anonymity guarantee or hope can be penetrated. But it's fundamentally very good, and it takes somebody with a huge amount of resources and money and will in order to sort of even then weakly deanonymize. And I think the most interesting thing that came out of the research is that it is much easier to prove an assumption about who someone is than it is just to collect data on everybody. So even in efforts to deanonymize, it's first necessary to make an assumption about who the endpoints belong to and then, with a great deal of effort, it's kind of possible to confirm the assumption. But even that is a weaker success on breaking anonymity.

Okay. With all that said, the Tor Project recently published, because they are a nonprofit, they published their tax exemption Form 990 which reveals a lot about the source of their funding. And the really nice bit of news here is that in 2017, which is what this recently filed form covers, the U.S. support had almost hit 50%. It was at 51%, with the balancing 49 being sourced, not from the public sector U.S. government, but from various private sector entities or other governments. The cofounder and the developer of the original software was a guy named Roger Dingledine. And he said...

Leo: Really?

Steve: Yeah, I know.

Leo: Poor guy.

Steve: Roger Dingledine. Sorry, Roger. Anyway, he said, in terms of percentages, while 2015 saw 85% of our funding coming from the U.S. government, 2016 saw the fraction drop to 76%, and 2017 we're down to 51%. So, I mean, even the fact that Roger's talking about this suggests that he, too, as a representative of the Tor Project, understands that people have always been made a little uncomfortable that the U.S. government is funding all of this. He said: "I should take a brief moment to explain how funding proposals work, for those who worry that governments come to us wanting to pay us to do something bad. There is never any point where someone comes to us and says, 'I'll pay you X to do Y.' The way it works is that we try to find groups with funding for the general area that we want to work on, and then we go to them with a specific plan for what we'd like to do and how much it would cost for us to do that, and if we're lucky they say okay."

So in 2017, breaking down the funding, was just shy of $800,000, $798,000, came from the U.S. government-backed Radio Free Asia; $635,500 came from the similarly U.S.-backed SRI International; $594,000 from the Swedish International Development Cooperation Agency, SIDA; $548,000 from the U.S. NSF, the National Science Foundation. And the State Department, which as we noted financed Tor's initial development and has sustained the project through its first years by covering most of its costs, has been reducing its involvement in Tor funding. In 2015 it was in for $200,000. It went up a little bit in 2016 to $218,000. But in 2017 it was down to $133,000.

So at the same time, last year the Tor Project raised a record $425,709 from its users, which is more than twice the funds it raised from users in 2016. But so far this year it's looking a little bit bleak at only $95,000 so far. As I think we've mentioned in the past, Mozilla has been a good partner with Tor. They've pledged matching funds for user donations. And overall, Mozilla independently of that has stepped up, boosting its contributions from $24,500 in 2016 to a whopping more than half a million, $522,188 last year. And they are expected to be a top contributor in 2018, as well. So yay for the

Mozilla Foundation for doing that. And DuckDuckGo also contributed $25,000 to the Tor Project in 2017, just because they wanted to support that.

It's also, Tor Project has also received in-kind service donations such as free cloud computing, free hosting, volunteer coding, translation services, and legal services, which they estimate for the purposes of their filing at around $806,000. So all told, the Tor Project's 2014 total revenue was $2.5 million. It grew to $3.3 in 2015; dropped by 0.1 to $3.2 million in 2016; hit an all-time high of $4.2 million last year in 2017. So Roger said that Tor's budget, even at the 2017 level, which seems like a lot at $4.2 million, remains modest, considering the number of people involved and the impact they are having.

He said: "It is dwarfed by the budgets that our adversaries are spending to make the world a more dangerous and less free place." So for what it's worth, their donation page is donate.torproject.org. And it's kind of a cool page. It's live. And if you watch it for a while, you might see it jump a bit. I went there and looked at it and happened to see the numbers change. So somebody had just donated while I was there. So it is really nice to know that a network designed not to provide anonymity, I mean, it's nice to know that in a network like the Internet, which was designed not to provide anonymity, such a facility like Tor exists and that it really does provide strong anonymity, which as we discussed previously can only be penetrated at great cost and with tremendous resources. So yay for those guys.

I mentioned an entirely foreseeable disaster at the top of the podcast. Believe it or not, a recently designed protocol for IoT devices uses UDP with no authentication. Which pretty much tells you all you need to know, at least our audience, about what is going to happen and in fact has started to. Leo, I know you jumped on the NTP protocol, the Network Time Protocol, for its use in amplifying DDoS attacks.

Leo: Right.

Steve: And of course we've previously talked about how DNS can be used. You make a simple query to a DNS server, and it returns a larger response. So if you spoof the source IP of your query, then the DNS server's response goes to the IP you spoofed rather than back to you. Same is true for NTP. In the case of DNS, it can be restricted so that only the users of the DNS server's local LAN have access. The same is true in the case of NTP. The problem is for protocols which are intended to be globally accessible and want to use UDP. Well, it turns out that there's a new one on the field. It's called C-O-A-P, CoAP, the Constrained Application Protocol. It is, self-described, it is at CoAP [C-O-A-P] dot technology, describes itself as a "specialized web transfer protocol for use with constrained nodes and constrained networks, meaning low-cost, low-power devices in the Internet of Things. The protocol is designed for machine-to-machine applications such as smart energy and building automation."

And as I said, here's the bad news: It supports the use of unauthenticated UDP packets in a simple query/reply mode, and a very small query can return a much larger, 50 times larger, reply. So it's a dream come true if it exists publicly on the Internet for DDoSers. It is essentially, in fact it can be thought of as a compactified HTTP. The protocol can support DTLS-style certificated-based endpoint authentication and encryption, like TLS over UDP. But when that's done, it's no longer tiny and simple, nor is it lightweight. So all that extra goodness has been eschewed in favor of just sending a query and returning a reply.

CoAP is new, having only recently been ratified into a standard four years ago in 2014. So we would expect uptake to follow the creation of libraries and then its manufacturing and deployment into devices. So what does Shodan have to say about it? After several initially quiet years, the number of publicly exposed, accessible, and answering CoAP devices has only in the last few months exploded, actually over the last year, exploded since a little over a year ago, beginning in November of 2017. It first appeared on the map in November of 2017, so 13 months ago, with about 6,500 devices. A month later, it was at 26,000 devices. Then by May of this year that number had grown to 278,000 devices. And today we're in the somewhere between 580 to 600,000-device range, according to Shodan.

Dennis Rand, who's the founder of eCrimeLabs, who has been monitoring the explosion of this protocol, watching this happen, believes that the reason for this explosion is CoAP's use as part of something known as the QLC Chain, formerly known as Qlink, which is a project aimed at building a decentralized blockchain-based mobile network using WiFi nodes throughout China. If you look at Shodan for a breakdown of the IPs where CoAP is available, China has 573, almost 574,000 of them. The next biggest, okay, 574,000, next biggest is the U.S. at less than 4,000, 3,831; Russia, 1,675; Canada, 327; Germany, 193. So these are virtually all in China.

And this recent rise in readily available and poorly secured CoAP clients, which is to say no security, has not gone unnoticed. Over the past few weeks, the first DDoS attacks carried out via CoAP have started to appear. ZDNet reported that a security researcher who deals with DDoS attacks, but who couldn't share his name due to employment agreements, stated that CoAP attacks have occurred occasionally during the past few months, and now with increasing frequency, reaching 55Gb on average. Remember, this is a factor of 50 amplification attack. So it is very simple to take compromised routers, bounce traffic off of them using their external UPnP exposure, bounce a little bit of traffic off of them from a spoofed IP, and have that traffic then hit a CoAP device, of which there are now 600,000, and have that device then amplify the traffic by a factor of 50 and send it on to its victim destination. So, yikes, 55Gb on average.

And one of the CoAP, the largest CoAP attack was seen at 320Gb; 55Gb, to put this in perspective, is an order of magnitude larger than the average size of a normal DDoS attack, which is about 4.6Gb, according to the DDoS mitigation firm Link11. Of these 580-some CoAP devices currently available on Shodan today, the same researcher told ZDNet that roughly 330,000 of them could be abused, so more than half, to relay and amplify DDoS attacks with an amplification factor of about 46, so roughly 50. And of course all this was foreseeable. It turns out the people who were doing this understood that UDP could be spoofed, but they did not build in a lightweight means of preventing that. What they unfortunately build in was, oh, gee, you could bring up DTLS.

And so there is a security protocol for CoAP, but nobody uses it because it requires - I saw it. It looked like maybe eight or nine roundtrips. And basically it's full weight DTLS, TLS on top of UDP protocol. Which is just sad. I mean, it takes no rocket science to figure out how to solve the problem. When I was writing the show notes I thought, well, okay, how would I solve it? One way would be to have the initiator first ask for permission by sending a small permission request packet to the endpoint it's intending to make a request from. The recipient would simply encrypt the requesting IP with its own locally unique private key and return that as a permit. It needs to retain no state. All it does is there's an incoming request for essentially permission to ask you questions.

When it powered up, it generated a random encryption key. While it's powered up, it uses that to encrypt any incoming permission requests of the requesting IP and sends that back as a permit. Then only requests containing the proper permit which match the requesting IP would be replied to. This would require one extra UDP roundtrip, very lightweight, much lighter weight than certificates, and that extra roundtrip would only be required once per endpoint pair, so the overhead would be minimal. The recipient never needs to retain state, and only the endpoint asking the questions needs to retain permits for however long it chooses to, and it's able to request a renewed permit if it has decided to discard it.

So again, it's not like these are hard problems to solve. There's a minimal implementation. This was very much like the stateless SYN solution that I came up with, which it turns out the industry had, Dan Bernstein had also come up with previously for Linux, which solved the problem of needing to maintain state when setting up a TCP stack. So again, this stuff is out there, but people are not bothering to take advantage of it. And as a consequence, now we have a rapidly growing number of new IoT endpoints that will happily multiply attack traffic by a factor of 46, and we're getting much larger DDoS attacks as a result.

So a little bit of, I guess this is errata because it corrects something that I just assumed was still true, but I'm very thankful for it. We talked last week about this debacle that Sennheiser has gotten themselves into, where they were found installing a single common root certificate in their HeadSetup application, very much the way Lenovo did back in the - what was it? I'm blanking on the name of the Lenovo mess.

Leo: Oh, yeah. Chatroom [crosstalk].

Steve: Somebody, yeah. Anyway, I got email from - so these were the guys at Secorvo who found the Sennheiser root CA. And I got email from Andre Domnick.

Leo: Superfish.

Steve: Superfish, yeah, the Lenovo Superfish debacle.

Leo: Thank you. [Vetman] wins. He was the first.

Steve: Thank you.

Leo: Then [Chickenhead].

Steve: So Andre Domnick of Secorvo send me a note saying: "Hey, Steve. Thanks for mentioning our small finding" - which nobody thinks was small. It was very cool. He says: "But just a short comment on the Chrome stuff. It is often mentioned in the media that Google Chrome is performing certificate pinning for the Google domains. But Google apparently has backed away from the idea. In our report you can find a picture of the current Google Chrome accepting the certificate. Additionally, we did not detect any unusual behavior of Chrome when presenting our cert to the browser."

So we can't take that as definitive, but it's interesting because I've been saying on the podcast every time it comes up that the Google browsers know, they have the fingerprints of the Google certificates. In fact, we know that it at least used to be true because this is how fraudulent certificates have been found, have been spotted in the wild in the past, was that Chrome said, hey, that's not from Google. So I just wanted to thank you, Andre, for the correction, and we'll see if we get any more information about that because maybe Google did, I mean, I've wondered how they could pin all their certificates because they're creating them like crazy and minting them on the fly. And it seems difficult to do that. But maybe they've backed away from that. I don't know. But I did want to mention that Andre says maybe we can't be counting on that.

Greg K. in Columbus, Ohio, regarding CA certificates, he says: "I'm fine with your rants about adding certs to trusted CA stores and middleboxes. However, I'm an advanced user. If I want to use my own internal root CA or my own middlebox to knowingly scan my own web traffic, I should be able to do that. I should not be restricted from adding whatever I want to my trusted cert stores on my network through group policy or whatever means. So I agree with your comments, but I am opposed to restricting the ability to do these things. I know what I'm doing, but others do not. So please be careful and responsible when you call for banning these activities to those individuals or organizations that can do it responsibly. I know you didn't exactly do that, but I would hope that you are being careful with your opinions."

And I'll just say, okay, I did mention that I also have my own root certificate. I've got one for this machine and one for localhost, which I discovered when I ran the Mark Russinovich or Sysinternals test against the root stores to make sure that nothing creepy had crawled into my root store. So Greg, I take your point. And I guess what I would opt for would be something like the dark screen, whoa, hold on, what you're about to do is installing a root certificate into your store. Are you sure you want to do this? That doesn't exist, a UAC-style "whoa" dialogue. That doesn't exist right now. And as a consequence, these things are just able to slip in without our knowledge. So I think it's necessary for some control to be provided. But I agree with you. An expert end user still needs to have the ability to do this.

Craig in Edinburgh said - his subject was "Self-signed cert surprise all the way from Azeroth." So I guess that's got to be a location in Edinburgh. [Reference is to WoW world.]

Leo: Edinburgh.

Steve: Edinburgh, oh, okay.

Leo: They don't say the "gh." You've got to say Edinburgh.

Steve: And he actually did say, I'm not kidding, after this, "No accent, please, Leo."

Leo: Oh, no. Well, I already blew it. I already blew it.

Steve: He said: "Given your audience, I bet I'm not alone in alerting you to this, but thought it was worth flagging up. After you mentioned Sigcheck" - that's the Sysinternals tool - "I ran it, expecting to be clean, as I don't install much, and was surprised to find one cert listed from Blizzard for Battle.net."

Leo: Uh-huh. Uh-huh.

Steve: He said: "Yes, I'm a longtime Warcraft addict. I could quit anytime I like, honest." He says: "I wondered what it was doing there and came across this post from a Blizzard employee." And I have the link in the show notes. I'll read it in a second. He says: "I don't think it's anything to really worry about as the last line says the desktop app generates a self-signed certificate that's unique to your machine and configures your system to trust it."

Leo: Uh-oh.

Steve: "So it's not the same as the Lenovo screw-up, but I had no idea it was there until today. It probably ought to be, at the very least, opt-in; don't you think?"

Okay. So Blizzard posted: "Our recent update to the Blizzard Battle.net desktop app made sure players could properly use features like logging into Battle.net via a social network, or joining a Blizzard group via an invite link. To facilitate these features, we updated the local web server to use a self-signed certificate to be consistent with current industry security standards.

"For those interested in more detail, using these features requires your web browser to communicate with the Blizzard Battle.net desktop app. Previously, the desktop app used a certificate signed by a public Certificate Authority, meaning that no modifications to your system certificates were necessary; however, this technique is incompatible with Certificate Authority policies, and we can no longer use it."

So, first of all, yay to Blizzard for not continuing to do something that is not safe because that would have required that their server running in the desktop had a private key for the certificate that they're using, and that's a big no-no.

Then they said: "While some browsers such as Chrome and Firefox are equipped to handle browser-to-app communication techniques, the changes were necessary for other browsers. For the time being, the desktop app generates a self-signed certificate that's unique to your machine and configures your system to trust it." And that is exactly the way it should be done. So again, bravo to Blizzard.

What unfortunately Superfish did and what, amazingly, Sennheiser repeated, was using a single certificate globally, such that in the case of back when Lenovo made this mistake with the Superfish software, which was an add-on, there was only the Superfish system installed a certificate system that was common to all Lenovo users who had that Superfish software, which then allowed spoofing of any site to all of those users. Same thing for Sennheiser. But a per-system certificate - and what I said last week, the mistake that Sennheiser made was not creating a per-system cert, which their installer could easily have done, which Blizzard did. So Blizzard did what they wanted to do correctly. So, yay. And Craig, thanks for sharing that experience and example.

Okay. So an important lesson and teachable moment, thanks to more troubles with Google+. David Thacker, VP, Product Management of G Suite, posted yesterday under the title "Expediting Changes to Google+." And I will quickly share what he wrote.

He said: "In October we announced that we'd be sunsetting the consumer" - the consumer, that's one key word - "version of Google+ and its APIs because of the significant challenges involved in maintaining a successful product that meets consumers' expectations, as well as the platform's low usage. We've recently determined that some users were impacted by a software update introduced in November that contained a bug affecting a Google+ API." And of course our listeners will remember that we covered a different exactly similar problem not that long ago.

He says: "We discovered this bug as part of our standard and ongoing testing procedures and fixed it within a week of it being introduced." So the point I'll get back to later is that they have something called "standard and ongoing testing procedures" which brought this to light. And that's significant. He says: "No third party compromised our systems, and we have no evidence that the app developers that inadvertently had access for six days were aware of it or misused it in any way.

"With the discovery of this bug, we have decided to expedite the shutdown of all Google+ APIs. This will occur within the next 90 days." Which essentially accelerates this by about half a year. He said: "In addition, we have also decided to accelerate the sunsetting of consumer Google+ from August 2019 to April 2019. While we recognize there are implications for developers, we want to ensure the protection of our users."

Under "details about the bug and our investigation," he wrote: "Our testing revealed that a Google+ API was not operating as intended. We fixed the bug promptly and began an investigation into the issue. Our investigation into the impact of the bug is ongoing, but here is what we have learned so far: We have confirmed that the bug impacted approximately 52.5 million users in connection with a Google+ API. With respect to this API, apps that requested permission to view profile information that a user had added to their Google+ profile like their name, email address, occupation, age, and so on were granted permission to view profile information about that user even when not set to public." So it was a glitch in this "should this be public or not," and the upshot was it published information that the user had put in, but had marked non-public.

"In addition, apps with access to a user's Google+ profile data also had access to the profile data that had been shared with the consenting user by another Google+ user, but that was not shared publicly." So again, another publication mistake. "The bug did not give developers access to information such as financial data, national identification numbers, passwords, or similar data typically used for fraud or identity theft. No third party compromised our systems, and we have no evidence that the developers who inadvertently had access to this for six days were aware of it or misused it. We have begun the process of notifying consumer users and enterprise customers that were impacted by this bug. Our investigation is ongoing as to any potential impact to other Google+ APIs."

"So," they said, "we will sunset all Google+ APIs in the next 90 days. Developers can expect" - so basically this is the last straw. They're like, nobody is using this. All it represents is the potential for exposure. We've already said we're going to shut this down. Nobody cares. We're just going to do it sooner. And probably after the initial announcement, and everyone just yawned, they said, oh, maybe we could have done it a lot sooner.

"So," they said, "developers can expect to hear more from us on this topic in the coming days and can stay informed by continuing to check the Google+ developer page." They said: "We have also decided to accelerate sunsetting consumer Google+, bringing it forward from August to April 2019. We want to give users ample opportunity to transition off of consumer Google+; and, over the coming months, we will continue to provide users with additional information, including ways they can safely and securely download and migrate their data." And they had a note in here which I'll skip to enterprise customers saying, don't worry, this is not going for you.

The importance to me of this phrase: "We discovered this bug as part of our standard and ongoing testing procedures and fixed it within a week of it being introduced." I want to formally congratulate Google for this, and I think it brings up an extremely and increasingly important issue. If a bug exists or is introduced at some point, it can subsequently be discovered, and history is showing us will eventually be discovered by any one or more of four possible entities.

It could be found by a malicious black hat hacker who leverages it to damage the company who created the software and/or their users, or customers who have trusted the company's offerings to be safe to use. As we've recently seen several times, it might also be discovered by a self-aggrandizing gray hat who does not themselves leverage the discovered vulnerability, but who discloses the bug's existence irresponsibly to both embarrass the company and endanger their users or customers until such time as it can get fixed.

And a perfect example are the three recent zero-days which we've seen tweeted. Or the bug might be discovered by a security researching white hat who responsibly discloses their discovery and quietly notifies the company to the problem so that it can be repaired. Or, finally, the bug might be discovered by the company's own employees, who can, as Google just has, immediately repair the oversight, determine the nature and extent of its impact on their users or customers' exposure, and notify them proactively of any possible consequences. While inadvertently outsourcing the discovery of unknown bugs to third-party security researchers is often what happens, it also often happens that those discoveries are made by those whose hats are not white.

So it's a gamble. And unfortunately it's a gamble that many companies make because they do not, cannot, or choose not to build in processes to find those bugs themselves. We know that those who create a system that's intended to be secure and bug free are the least able to attack it and find fault in their own work. I mean, all of our experience demonstrates that. So some kind of internal red team facility needs to be created. It could be a separate group within the company who pride themselves on punching holes in the company's own offerings, or an automated system which continually scans, fuzzes, and searches for defects and odd behavior.

The point is the establishment and maintenance of any such proactive vulnerability discovery system is not free, and the value it returns to the company can be difficult to justify. One of the lessons we have ample opportunity to learn and relearn on this podcast is that security is not free. And companies do not make headlines for not having serious security incidents. They only make headlines when their security fails. So I want to salute Google for clearly having such processes in place, which are working, and often find problems themselves rather than relying even inadvertently for them to be discovered by others. We learned from Paul and Mary Jo that Microsoft had much more than ample feedback from users of their most recent Windows 10 feature update disaster, which was warning of every problem that later surfaced, at great cost and embarrassment to them and their users.

So in this instance Microsoft had such a system, but it was there apparently pro forma, just in name. They weren't actually looking at any of the feedback and the data that it was producing. So you also have to heed what the system you have put in place is producing. And we also know that they understand that now. Hopefully their behavior will change. And I was reminded of this way back 30 years ago when SpinRite was young, and I had an entire bullpen of technical support people on the phone with customers. We learned that there's no such thing as a one-off. Anything that happens to one user will happen to others. But that's a lesson that takes some time to sink in. And it's not good news.

So some people in some organizations prefer not to believe it. They think, oh, well, we don't know why that happened. It probably won't happen to anybody else. No. Given enough time and enough people, it will. So you have to pay attention to every one of them. So anyway, props to Google. In this regard, I think they are a company worthy of emulation. And I really do hope that, as the technical press continues to note the scope and the cost of these breaches, that companies understand that it is possible for them to find them themselves, rather than to rely on outsourcing it to, in many cases, black hats who can also hurt their users. So there.

Leo: So there. Is that it?

Steve: That's it.

Leo: That's all we're going to say? You took me by surprise. Thank you, my friend. Steve Gibson comes to us courtesy of the Gibson Research Corporation, providers of the finest hard drive recovery and maintenance utility known to man, a.k.a. SpinRite. Go to GRC.com. Get your copy. And while you're there you might check out all the other great stuff Steve offers, including this show - not only audio of the show, but also transcripts. It's the only place you can get those, and they're searchable, which makes it really easy to go through all 693 episodes, find the topic you're interested in, going back all the way to Honey Monkeys, honeypots and all of that, 13 years ago.

You can also get audio and video from our site, TWiT.tv/sn. And then I think the best thing to do is start your collection by subscribing to Security Now!, getting every episode the minute it's available on a Tuesday evening. You can collect all 693, or 694 next week. I will be back, so will Steve, as we are most Tuesdays about 1:30 Pacific. We were late today, I apologize, but usually around 1:30 Pacific. That's 4:30 Eastern, 21:30 UTC. If you want to watch live, go to TWiT.tv/live. Chat with us in the chatroom at irc.twit.tv. That's all the people watching live. So that way you can, when Steve and I say, "What was the name of that [Lenovo] rootkit?" you can say, "Superfish, Superfish, Superfish." Patrick Delahanty says: "Collect 'em and trade 'em with friends." Got to get them all. Steve, we will see you next week. And I thank you so much for another great show.

Steve: Thank you, buddy. That will be the last podcast of the year since we have two weeks off, since Tuesday is when Christmas Day and New Year's Day fall. So I'll try not to forget how to do them in that time.

Leo: I doubt you will. I might.

Steve: I was joking with Lorrie. She said the same. She says, "I don't think you're going to forget."

Leo: Thanks, Steve.

Steve: Okay.

Leo: See you next time.

Steve: Right-o. Bye.