Russian Meddling Technology

Description: This week we examine and discuss the appearance of new forms of Meltdown and Spectre attacks, the legal response against Intel, the adoption of new cybersecurity responsibility in New York, some more on Salon and authorized crypto mining, more on software cheating auto emissions, a newly revealed instance of highly profitable mal-mining, checking in on Let's Encrypt's steady growth, the first crack of Windows uncrackable UWP system, Apple's wacky Telugu Unicode attacks, a frightening EternalBlue experiment, another aspect of crypto mining annoyance, a note now that Chrome's new advertising controls are in place, and a bit of closing-the-loop with our listeners.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-651.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-651-lq.mp3

Then we conclude with a look into the technology that was revealed in last week's indictment of election-meddling Russians and, from a practical technology standpoint, the feasibility of anything changing.

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We've got more news about cybersecurity, of course, including, yes, more Spectre and Meltdown news. Oh, man. Plus how the Russians hacked Twitter and Facebook, and a bitcoin miner that nearly brought T-Mobile to its knees. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 651, recorded Tuesday, February 20th, 2018: Russian Meddling Technology.

It's time for Security Now!, the show where we cover your security online with this guy right here, Mr. Steven Gibson from GRC.com. Live long and prosper, Steve. Hello.

Steve Gibson: Yo, Leo. Great to be with you again, as always. So we plow in, well, past Episode 650. I don't know why that number seems so significant to me, but we're on 651 today. And a piece of information became available last Friday that I thought was interesting. And also Techdirt's response, Mike over at Techdirt, sort of a little bit of a reality check as a consequence. What I'm talking about is what we learned from the Mueller investigation about the technology that Russia allegedly used for messing with social media. And I don't care about the politics. What was interesting to me was what

they did in order to pull this off.

And this is relevant because right now, sort of in the wake of that, there's all this comeuppance, all this furor about the responsibility of Google and Twitter and Facebook and Instagram to have dealt with this or to have known better, and what can we do coming up. And of course this put me in mind of the FBI's arguing that encryption was bad, and that we needed a golden key. It's like, well, there are things that technology can do, and there are things that technology can't.

So now that we know a little bit more about how this was done, we have something to talk about from a technology standpoint which I think is sort of interesting about just the nature of the Internet and how little control over it we really have. But lots of news this week, whereas last week, actually the last couple weeks have been kind of quiet. So I want to examine and discuss the appearance of, unfortunately, new forms of Meltdown and Spectre attacks.

Leo: Oh, boy.

Steve: Uh-huh. There's just some beautiful research out of Princeton that revealed something that I guess we should have suspected or expected. Also the legal response against Intel I want to touch on briefly. They had to file their annual 10-K report, so we know we've got a little glimpse into what would otherwise be proprietary about what the Meltdown and Spectre vulnerabilities have done in terms of backlash, which I think is unfortunate because I don't think this is Intel's fault, any more than it's - it was the nature of the technology.

Also there's some new cybersecurity legislation which is going into effect, actually last week and also in coming weeks in New York, which I want to cover, and we'll talk about it. A little bit more on Salon and authorized crypto mining. We have a little update on software that cheats auto emissions. Oh, and also maybe a new coined term and a newly revealed instance of highly profitable, what I would call "mal-mining."

Leo: There's a word I don't want to hear again.

Steve: Mal-mining, yes.

Leo: Mal-mining.

Steve: I want to check in on Let's Encrypt's steady growth, which as you noted is the graph that I chose for the Picture of the Week. It just keeps on trucking. We also have the first crack reported of Windows', now in quotes, "uncrackable" UWP system. I wanted to touch briefly on the iOS and across-the-board updates we got and the crazy Unicode attacks that beset Apple. A frightening EternalBlue experiment which was contracted for by a company wanting to understand the nature of their vulnerability. And everybody was a little surprised by what it revealed. Also we have another new aspect of crypto mining annoyance.

I wanted to make a note now that Chrome's new advertising controls are in place that we've talked about previously, that I know you talked about over the weekend. A bit of

closing the loop with our listeners. And then, as I said, I want to look at some of the technological cleverness, unfortunately, that apparently Russia got up to in 2016 and whether there's actually anything that we can do about it. So I think a chock-full podcast.

Leo: Nice. Very excited.

Steve: From our "attacks never get worse, they only ever get better" department, it was inevitable that there would be additional forward motion on the whole Meltdown and Spectre side. I loved the proud Princeton professor's tweet which started this. She tweeted: "My PhD student @carolinetrippel developed research tools/techniques to synthesize security exploits. Yes, her tools find #Spectre and #Meltdown. But they've also discovered two new related-but-distinct vulnerabilities using cache coherence protocols." So that's from the professor of the PhD student, or candidate, I guess, who generated a beautiful paper.

The research paper, it was done along with Nvidia, is titled "MeltdownPrime and SpectrePrime: Automatically Synthesized Attacks Exploiting Invalidation-Based Coherence Protocols." Which is sort of the academic formal way of talking about the idea that you have some situation where caching is valid or invalid, and you need to maintain so-called "cache coherence," that is, you need to make sure that what's in the cache is coherent with what's stored in main memory so that the cache is a faster access scratchpad, but you need to maintain coherence. And that's an issue when you've got multiple things accessing a shared cache. You need to make sure that they both have the same view of the data at all times.

Anyway, so these researchers developed a tool using mathematical rigor to formalize essentially the underlying architectural characteristics which enabled the creation of the Meltdown and Spectre side-channel attacks. So they sort of stepped back and generalized the whole problem, saying that, okay, wait. So here's a couple, you know, Spectre and Meltdown are specific instances of an attack. But is there a broader problem here? And what they answered, you know, they came up with, oops, yes.

And one has to think that, if not before now, then certainly since the summer when this was first revealed to them, revealed to Intel, sorry, revealed to Intel, that Intel would now have similar tools at their disposal, I mean, one would hope; although it's clear that they didn't previously, or they didn't think that this could be leveraged into an attack, you know, it's sort of unclear how much of this came as a surprise because, as we talked about when we originally talked about the Meltdown and Spectre problems, there was some awareness of this decades ago, that this was a potential problem, that is, that there was a way to leverage speculation and caching side channels.

So anyway, this new tool that they built allowed the researchers to apply mathematical rigor to develop new software-based attacks from the description of the CPUs microarchitecture, and then this tool essentially writes the code which is then able to perpetrate the attack. It is specific to specific microarchitectures. On the other hand, what we've seen is that, across a large set of microarchitectures, Intel has pretty much done the same thing, which is why they're in trouble across so many years of their successive chips.

So anyway, in their abstract, to give you sort of a sense for the rigor of this, the abstract, just the beginning of it reads: "The recent Meltdown and Spectre attacks highlight the importance of automated verification techniques for identifying hardware security vulnerabilities. We," they write, "have developed a tool for automatically synthesizing

microarchitecture-specific programs capable of producing any user-specified hardware execution pattern of interest. Our tool takes two inputs: a formal description of a microarchitecture in a domain-specific language" - meaning that they invented their own language to describe hardware microarchitectures - "and a formal description of a microarchitectural execution pattern of interest, for example, a threat pattern." Then they write: "All programs synthesized by our tool are capable of producing the specified execution pattern on the supplied microarchitecture."

So they succeeded. And when they applied this tool to the Intel microarchitectures, they found previously unsuspected additional problems, which is the real takeaway from this. Thus they call their paper "MeltdownPrime and SpectrePrime." Now, in some of the press's coverage of this, there was the point raised or the question raised whether this meant that Intel had to go back to the drawing boards yet again with firmware updates. And it's not clear one way or the other whether that's going to be the case, whether the existing firmware fixes will fix this.

But what we do have as a consequence of this very rigorous academic research of architectures is first a nice step forward in the science of the security of our hardware platforms, which as I said would have been nice if we had it quite some while ago, but at least we have it now and will moving forward, which may be the way that we're able to get the best of both worlds. We are not wanting to sacrifice undue performance that we've been taking for granted and incorporating into everything we do for the sake of security. But neither do we want architectures that are exploitable.

So I imagine there's a lot of head-scratching going on at Intel and in other chip manufacturers, recognizing that this is, as we said at the beginning of the year and spent weeks of this podcast talking about, this was not a small event in the history of computer science and this industry. This was a big deal, and the repercussions are still being felt. And this kind of fundamental research is just - is great to see. So props to these guys for pulling that together.

And in a kind of a coincidence, because of the timing of Intel's annual report filing, there's a - it's known as the 10-K form that has to be filed with the U.S. Securities & Exchange Commission, our U.S. SEC, which publicly traded companies have to file, where they formally lay out like where they stand, what the known risks to purchases of shares of their corporation are and so forth. It's a long document. It's a 201-page PDF. But coincidentally, page 124 of that PDF has to and did address the consequences of this Spectre and Meltdown on Intel. We talked about it at the time that the Intel stock suffered as a consequence when people like, oh, my god, this is a big deal.

So as a consequence of this report filing, we learned a few things. They said - and I'll give a little bit of intro here just to set the stage. In June, and so this is Intel now in their formal 10-K filing: "In June 2017, a Google research team notified us and other companies that it had identified security vulnerabilities - now commonly referred to as Spectre and Meltdown - that affect many types of microprocessors, including our products. As is standard when findings like these are presented, we worked together with other companies in the industry to verify the research and develop and validate software and firmware updates for impacted technologies.

"On January 3, 2018, information on the security vulnerabilities was publicly reported, before software and firmware updates to address the vulnerabilities were made widely available. Numerous lawsuits," they wrote, "have been filed against Intel and, in certain cases, our executives and directors, in U.S. federal and state courts and in certain courts in other countries relating to the Spectre and Meltdown security vulnerabilities." So this they have to disclose per the requirements of the 10-K filing with the SEC.

They wrote: "As of February 15, 2018, 30 customer class action lawsuits and two securities class action lawsuits have been filed. The customer class action plaintiffs, who purport to represent various classes of end users of our products, generally claim to have been harmed by Intel's actions and/or omissions in connection with the security vulnerabilities and assert a variety of common law and statutory claims seeking monetary damages and equitable relief. The securities class action plaintiffs, who purport to represent classes of acquirers of Intel stock between July 27, 2017 and January 4, 2018" - so that period after which Intel knew about this and the world found out about it - "generally allege that Intel and certain officers violated securities laws by making statements about Intel's products and internal controls that were revealed to be false or misleading by the disclosure of the security vulnerabilities. Additional lawsuits and claims may be asserted on behalf of customers and shareholders seeking monetary damages or other related relief.

"We dispute the claims described above and intend to defend the lawsuits vigorously. Given the procedural posture and nature of these cases, including that the proceedings are in the early stages, that alleged damages have not been specified, that uncertainty exists as to the likelihood of a class or classes being certified or the ultimate size of any class or classes if certified, and that there are significant factual and legal issues to be resolved, we are unable," Intel concludes, "to make a reasonable estimate of the potential loss or range of losses, if any, that might arise from these matters."

So I regard this as unfortunate. That is, I'm sure this is business as usual. I'm sure that for a company Intel's size there are always people who are upset. They buy stock. It goes down. They sue people and try to find some relief. We know that class actions themselves can often be kind of sketchy, where the individual plaintiffs haven't suffered damage, but they're aggrieved or feel aggrieved. Who knows how this is going to come out? I'm just sort of sad that something which from a technology standpoint you could argue everybody was doing, and we were all happy, we were all benefiting from the fact that we were using a chip architecture which could be cleverly leveraged to create some information leakage, but in return we were getting performance that we're now being asked to give up to varying degrees.

And as I'm keeping an eye on this, we're still as an industry learning more about what this means. In some cases - and I'm kind of watching Intel and AMD and ARM. And in some cases the documentation is being changed in order to make some of the instructions more clear that can be used for flushing, speculation, and synchronizing multiple cores. In some cases the behavior of the instructions is being subtly changed in order to create more rigor where some was lacking. But I'm sure Intel is adequately able to defend themselves and will be able to bring out experts that say, yeah, this was unfortunate, but this wasn't negligent. It's hard to see that that was.

And also it's certainly the case that delay from Intel's learning about this to disclosing was also reasonable because they had a lot of work to do, as I've said in the last few weeks, in order to remediate the consequences of what Google found and reported. And as we also know, they weren't quite done by the time this all went public, and everybody kind of tripped over themselves when that finally happened. So anyway, this is a huge, huge event for our industry, and we're still seeing repercussions from it.

Also last week, later in the week, New York State's new cybersecurity laws, which were actually put on the books in March of 2017, but the nature of them were that at some given date certain the following things had to happen by. That happened, the beginning of that happened last week. So what has happened is there's essentially an effort to create more accountability about the cybersecurity efforts that major corporations, largely banks, but also insurers and financial service organizations, have put in place in

order to protect their customers. So it affects over 3,000 banks.

And essentially the legislation requires that senior executives at the Chairman of the Board of Directors or a senior officer like a CEO personally certify that their computer networks are protected by a cybersecurity program appropriate to their organization's risk profile. This is from the New York Department of Financial Services, which are attempting to create some accountability and oversight so that, I mean, and we know that making these sorts of statements doesn't prove anything. It's different than audits from outside firms and active enforcement. But it feels like it's a start.

As of last week, these more than 3,000 firms had to have a very top-level executive formally attest to the following six points: that their firm had adopted a cybersecurity program appropriate to their company's, often a bank's, risk profile; that they had adopted cybersecurity policies designed to protect the bank's information systems and the customer data they hold; appointed a chief information security officer responsible for overseeing and implementing the bank's cybersecurity program and enforcing its cybersecurity policy, so somebody to deal with those first two points; engaged, qualified cybersecurity personnel, either staff or contractors, to work with that company's CISO in managing the company's risk; developed an incident response plan; and taken steps to control privilege access to its IT network.

So those all seem like reasonable things to ask for a large organization which is doing all these things. But what we have found in past years when specific events have happened, and we've covered it on the podcast, it's like, wait, you don't have a cybersecurity plan? You don't have any staff who's doing this? I mean, so what we've found is that things have just been negligently lacking. So at least this is intended to focus the attention of corporate management on the need for these things. And they had plenty of time to pull this together. So this isn't external audits on an ongoing basis, but it's a start.

Then two weeks from now on the beginning of March, on March 1st, is the second implementation deadline for this same legislation, which will then be a year old, and five more much more substantial regulations and provisions come into force. They have to have implemented either continuous monitoring or periodic penetration testing and vulnerability assessment for their network, so active participation in this; conducted a full-scale risk assessment of their information systems which will be used to inform their cybersecurity program; implement multifactor authentication for remote access, and perhaps more widely used as indicated by what the risk assessment showed; provide regular cybersecurity awareness training for their staff; and begin annual reporting from the CISO to the board of directors.

So again, things that you would hope any large, especially a large banking organization, where this stuff is critical, would already have well in place. But what we found is no one was really looking until we found instance after instance where this wasn't done. So New York State has said, okay, we're going to make the management of these major corporations attest to the fact that these things have been done. So a step forward, if nothing else.

We talked last week about Salon hoping to monetize opt-in cryptocurrency mining. And I read, as a consequence of the news of that, a number of interviews of Jordan Hoffner, who is the CEO of the Salon Media Group. And there wasn't a lot of information there, although I did get some that we'll talk about in a second about this notion of authorized mining. The most curious, to me, and disingenuous thing that I saw from these interviews is this question about CPU usage percentage because in being interviewed the interviewers were saying, you know, in one case I remember one interview said, well, you know, my CPU went to 90% of usage while it was on the FAQ page after I gave it

permission to mine.

Oh, and by the way, I meant to say that I reported last week that I had read reporting that stated that the mining started when you acknowledged that first interstitial blocking page, but before you gave explicit permission. It turns out that was true, but then stopped being true. So they changed the behavior of their code, which may be why, Leo, when you tried it, it didn't just take off and pin your CPU. So they cleaned that up a little bit, and that was good.

But the problem is people are saying, whoa, 90% CPU usage. And then this Jordan, and I don't know how technical he is, but he was saying, "Oh, yes, that's one of the things that we're looking into is there seems to be a great disparity in how much CPU is used. Sometimes it's 5%, sometimes 3, sometimes 90." And I'm thinking, okay, well, first of all, it's probably the case that you have something like uBlock Origin in place. It happens that uBlock Origin blocks any access to Coinhive from where those scripts are generated. So it's going to prevent essentially any CPU usage of mining.

But as I mentioned before, my real concern, I mean, I'm sort of bullish about this concept, that is, I can sort of see there's something here that I imagine is going to - we're going to see some traction on this. I think it's going to take hold. But JavaScript is the wrong way to do this. It's wrong because it is atrociously low yield. So as a consequence it doesn't work for anybody. It means that for value to be generated, a crazy amount of CPU resource, like yes, 90, 95, 100%, needs to be taken. But as a consequence of the inefficiency of JavaScript mining, which is hugely, as we know, it's crypto-intensive, the amount of dollar value or cryptocurrency generated for this in return for massive CPU use is minuscule. So it's going to be incredibly wasteful as system resources.

And so what we need for this to succeed is for mining to move into the browser, like to be a native feature of web browsers, if this is what we want to have happen, where for example it's able to access the system's GPU and perform reasonable yield mining with all of the proper controls and metrics and so forth. So the concern is, my concern is that it feels to me like there's a nugget of something here. Also, users are going to have to understand that, yes, if they don't want ads, they're going to have to provide value, and the value will be shown as their system generating more heat, their CPU suddenly jumping to just shy of 100. You know, you still want the system to run, but essentially mining will take up all of the other available resources. I mean, that's the tradeoff.

So anyway, it's just been sort of interesting, I mean, there's a lot of interest in this. And I think it makes sense. The problem is I hope this doesn't stub its toe because we're sort of premature in having a means for it being cost effective, if it turns out it is. There's just no way that doing this with JavaScript is cost effective. It immediately needs to get moved natively into the browser. Or, I mean, as I was imagining this, you could imagine installing a browser-mediated agent on the user's computer that would work, too. But it just makes more sense if it's native to the browser. The browser just, you know, which is an app running on the system, just has the ability, with permission, to suck up as much CPU as you choose to give it. At this point, it's doing that, but horribly insecure.

Now, what's really interesting is that I've also read a lot about these Coinhive guys. And I've seen my comment on this podcast retweeted a number of times, where I said several months ago when this was beginning that it was impossible to not at least partially hold the Coinhive people responsible, that they had to be complicit in this because, in the early instances, all of the instances of the use of their script seemed malicious, I mean, seemed to be, you know, it was all, to use my new term, mal-mining. So they've spun off a subsidiary called AuthedMine.com, A-U-T-H-E-D-M-I-N-E dot com.

And it's authorized mining by Coinhive.

And so if you go to AuthedMine.com - and, by the way, I was told that's what Salon is now using - you get the following: "A Note to Adblock and Antivirus Vendors. There is no need" - so this is an open letter. "There is no need to block AuthedMine.com." Now, you can imagine, they're asking people not to because Coinhive has gotten immediately blacklisted across the industry. "There's no need to block," they write, "AuthedMine.com or any scripts hosted on this domain. AuthedMine" - god, that's hard to say - AuthedMine.com offers a Monero miner that can be embedded into other websites. This miner will only ever run after an explicit opt-in from the user. The miner never starts without this opt-in."

Now, I was initially skeptical, but I kept reading, and I got convinced. They said: "We implemented a secure token to enforce this opt-in on our servers. It is not circumventable by any means, and we pledge that it will stay this way. The opt-in token is only valid for the current browser session (at max 24 hours) and the current domain. The user will need to opt-in again in the next session or on a different domain. The opt-in notice is hosted on our servers and cannot be changed by website owners. There's no way a website owner can start mining without the user knowing." And then they say, "Click here to see how the Opt-In looks." They say, "A detailed and technical explanation of the Opt-In can be found in our documentation. We believe, they conclude, that browser-based mining can be a viable alternative for intrusive and annoying ads if used honestly and with consent by the user. We kindly ask adblock and antivirus vendors to support us. Please help us build a better web. Cheers." Signed, Coinhive.

So this is a, from all appearances, a well-designed, properly structured and architected solution that's still script-based, and it'll be interesting to find out how much money, how much value can be mined from JavaScript. Certainly, it would be great if it's enough to be viable now. Then we have a solution that could potentially work today. If not, then we know that we can increase the leverage.

So then over on Coinhive they said of their AuthedMine branch, "A Non-Adblocked Miner." They wrote: "Shortly after the launch of Coinhive, several adblockers have begun blocking our miner. This is unfortunate because we intended Coinhive to be an alternative to ads, precisely for users with adblockers." Eh, okay. Maybe there's a little bit of history being rewritten there. Of course we have no idea what they intended.

"However, we have to acknowledge that the decision to block Coinhive was understandable as it was possible to run the miner on a webpage without asking the visitor for consent or even informing them. Even some antiviruses now consider our JavaScript miner as a threat, which makes it difficult for website owners to use Coinhive at all. We implemented AuthedMine as a solution to these problems.

"The JavaScript Miner, Simple UI, and CAPTCHA" - so they have three services, JavaScript Miner, Simple UI, and CAPTCHA - "when loaded from AuthedMine.com will never start without asking for consent from the user or" - and they say from for the Simple UI and CAPTCHA - "letting them explicitly start mining through a click. We realize this opt-in may be clunky and not fit all too well with your use case, but we strongly believe that being honest with the user will ultimately be beneficial for users and website owners alike. Neither the JavaScript files on AuthedMine.com nor the domain names are currently blocked by any adblockers or antivirus. We will talk to adblock and antivirus vendors so it will hopefully stay this way."

So I think this is clever and entirely appropriate. And it does, like right now, give websites that want to experiment with this a useful, valid, non-blocked, I mean, truly

opt-in, authenticated way of using browser-side mining in order to generate revenue for themselves. So I think this is very interesting. I think Coinhive did the right thing. It's clear that the Coinhive domain got immediately blacklisted. They're being very open. They've got a technology that makes sense.

I understood what they meant when they said we understand it may interfere with the flow of a website. You can imagine that a website would just like to get - it would itself like to get permission from the user, and then have the mining proceed, rather than require another opt-in from this AuthedMine.com that the website doesn't control. But there's no way to do that without opening yourself, as Coinhive did, to malicious use; to, well, we'll call it mal-mining. So anyway, it'll be interesting to see how this proceeds and whether there is uptake in adoption of this. In return for people's CPUs being pinned, revenue gets generated. So I think this is exciting. We'll see what happens.

Leo: On we go. Onward.

Steve: Yeah. So following up on my grumbling about the inefficiency of using JavaScript for mining, this is a little out of order in my show notes. I'm just telling Elaine that since she uses the show notes to do the transcribing. But of course we were talking last week about BrowseAloud, which was the site whose servers were breached, and the mal-mining Monero Coinhive was injected into, remember, 4,275 websites. Many of them were governmental. CUNY was among them, as was USCourts.gov and a bunch of others.

Well, we don't know who the creep was that did this, but as a consequence of the way crypto mining works, we know what success or not that particular crypto mining address had over the course of that campaign. And in something of a funny quip, I saw someone comment that it was noted that browser-based crypto mining relies upon placing the mining code on sites people actually visit. So again, a lot of sites got infected. Many of them may have been low traffic. But one way or the other, JavaScript-based mining, and this whole campaign infecting 4,275 websites, generated, minted a grand total of $24.

Leo: Oh. Oh.

Steve: So, yeah. I don't think anybody won from that proposition. Either those were not great sites to inject the mining code onto; or, as I said, it's just not generating a lot of money. And if it's going to be viable, we need to fix that.

So I did want to touch on this question that we first addressed when VW got caught playing games with the emissions control software on I believe it was the diesel VWs. That caused sort of a spotlight to be cast upon other potential violators. And it was reported in a German newspaper just this last Sunday that U.S. investigators, in looking at Mercedes, have found that its cars appear to have been playing similar games.

This German newspaper wrote that the documents which they - and these were private documents that the newspaper got a hold of - showed that U.S. investigators had found several software functions that helped the Daimler cars pass emissions tests, including one which switched off emissions cleaning after 26 kilometers of driving. They found another function in the software which allowed the emissions cleaning system to recognize whether the car was being tested based on speed or acceleration patterns, and they also found emails from the engineers involved, questioning whether these software

functions were legal. So it looks like maybe more than just VW was playing games, and that it's sort of a consequence of how complex our cars have become and the fact that companies just believe they're able to tuck away whatever technology they want to out of sight.

Now, the flipside of the inefficiency of JavaScript mining is the efficiency of mining natively on high-performance machines like big iron Internet servers. It turns out that the Israel firm Check Point, who is often bringing us interesting bits of news, announced last Friday that they had uncovered the footprint of a large hacking operation targeting something known as "Jenkins servers" which were left connected to the Internet. These Jenkins servers are Java-based technology which are used for developers staging their software. It turns out that they're intended to be used internally, but there were on the order of 25,000 of them exposed on the Internet.

Okay. So as always, there was a well-known patched vulnerability. This was part of a CVE that was issued last April 26th, so 10 months ago, that allowed unauthorized and unauthenticated remote code execution on these Jenkins servers, using something known as a Java serialization/deserialization flaw. Java uses a lot of data structures, and so Java objects are very structure heavy. The serialization in Java refers to a technique for turning one of these structured objects into a byte sequence which then allows it to be transmitted or stored, and deserialization is the reverse process. The deserializer reads the sequence of bytes and reassembles the original JavaScript object.

Well, it probably won't be any surprise to listeners of this podcast that the deserializer is an interpreter, that is, it is an interpreter of the serialized sequence. And as we often talk about, interpreters are very difficult to get right. They are a continuing source of security problems. And it turns out that there was a high-profile deserialization problem in Java which afflicted these Jenkins servers. It was known about 10 months ago. It was patched. But these servers were not updated.

So despite having been fixed over a year ago, the attackers were able to leverage this known remote code execution vulnerability in more than 25,000 exposed servers. And much as we're able again not to identify who it was who obtained the money, but we can tell from the block chain technology, the cryptocurrency mining technology, how successful they were. This had been active for months and allowed them to already, I mean, it's active right now, as of this reporting, which was just Friday, they have already cashed out more than 10,800 Monero, which is over $3.4 million.

So here's an example of a crypto mal-mining operation which is succeeding, is running on these Jenkins servers, as a consequence of the fact that it is running on the bare metal. It's not JavaScript. It's actually running natively on probably very powerful machines. I would imagine that these Jenkins servers are limping along because all of their CPU resource is being sucked up by a crypto mal-mining operation that doesn't care about sparing any CPU cycles for what the machine was supposed to be doing. And as a consequence, whoever this is has made $3.4 million USD, which is some serious coin. Thus the strength of the impetus to do cryptocurrency mining.

And it's worth noting that, with remote code execution vulnerability such as these servers have, bad guys could have gotten up to all sorts of much more serious, in terms of security penetration, corporate espionage, and digging around in people's network mischief. But what do they choose to do? They choose to combine cryptocurrency with their access.

BleepingComputer, a terrific site that we refer to often, had a terrific summary reporting on this. In their coverage, where they quoted over 2,500 Jenkins servers left exposed

online, they said: "Attackers aren't the only ones who've noticed the large number of Jenkins servers available online. In mid-January, security researcher Mikail Tunc published research highlighting that there were over 25,000 Jenkins servers left exposed to Internet connections at the time of his research. Also on Friday" - that is, last week - "FireEye released new research on other hackers leveraging a different flaw which is being used to infect Oracle WebLogic servers with malware. This vulnerability has been under active exploitation since early December 2017, and one group has already made more than $226,000." So there's a different cryptocurrency mining operation running on Oracle WebLogic servers.

Then Bleeping Computer says: "Besides Jenkins and Oracle WebLogic servers, hackers are also targeting Ruby on Rails, PHP, and IIS [Microsoft] servers, also deploying Monero mining malware. Trend Micro fears," they write, "that two recently disclosed CouchDB vulnerabilities will also soon be exploited in the same way." They write: "Monero mining malware is already this year's biggest malware trend and problem, with numerous malware distribution campaigns spreading such payloads on any unsecured computer and server that crooks can get their hands on."

That's what Bleeping Computer said, and then I was reminded of the famous quote, the Willie Sutton quote, thinking that perhaps we should be referring to all these as Willie Sutton attacks, going where the money is. So it's interesting that, with the advent of cryptocurrency and exchanges that essentially take the cryptocurrency from being virtual to being real. As we've been talking about recently, there's now a great deal of pressure to get mining malware running on computers wherever you can find them. And it's generating dollars, not just random mischief or worms that propagate for their own entertainment, but for no other purpose. Now, unfortunately, for better or worse, we've given them purpose.

I just wanted to make a quick note, this is from the "some secrets cannot be kept" department, that reports of the first "uncrackable," in quotes, Windows 10 Universal Windows Platform, that's the UWP, based game has reportedly been liberated from its captivity. Who knows why they chose this one. The game is Zoo Tycoon Ultimate Animal Collection which has been cracked. The cracking group CODEX says that in order to do this they needed to successfully penetrate five separate layers of DRM protection.

So doing this was not easy. But from the beginning of this podcast we've noted that there are some things which cannot be protected. Our first example was the encryption of DVDs back when there was that effort underway because, if the console in your living room needed to decrypt the DVD to show it to you, well, everything you needed to know to decrypt it was right there. Similarly, if someone is playing a game on their computer, despite any efforts that anyone could make in order to corral it and control it and prevent it from getting loose, it's running on a machine which is able to penetrate all those five layers. So it's just a matter of someone patiently doing the reverse engineering in order to pry that apart.

We all got updates, those of us using Apple stuff, iOS for iPhone and iPad, tvOS for Apple TV, watchOS, and even the macOS for a problem that Apple clearly knew about before it went public. We know that Apple knew about it because the planned next sort of semi-major release of iOS, that's 11.3 that we should be getting soon, that's the one that I'm looking forward to because it removes the speed throttling of older devices, or at least gives a user some control over it and notification, and also apparently clear reports on the health of your battery as the phone sees it. That's the one that's coming, 11.3. Apple it turns out had a problem with 11.2.5. It's known as - is it pronounced "Telugu"?

Leo: You know, I don't know. We were debating it. I think it's Telugu, something like that. It's an Indian dialect.

Steve: Yeah, yeah. It's the third most spoken language native to South India. And it turns out that - and also there was another, was it - I don't think it was Belize. Bulgarian? There was some other language character that was later discovered. There's a fabulous posting, if anyone is interested in, like, just for grins, understanding exactly every technical detail of what this crash was about. I've got a link to a page on GitHub.io, a blog page there that really dissects this problem, courtesy of Simon Zerafa, who sent that to me.

Anyway, the bottom line is that sending a Unicode character sequence using these Telugu characters was able to crash these various devices. Apple knew about it. It doesn't, for example, crash the beta of 11.3, so they fixed it there. But news escaped. It began getting exploited in the wild, and so Apple was forced to push out a quick update, basically that they'd already fixed and would have disappeared without anyone knowing about it in 11.3. They had to move us.

Leo: It's a Bengali dialect. It's another Indian dialect, yeah, or Indic.

Steve: Oh, Bengali, that's right. Thank you. I knew it was a "B" word.

Leo: Yeah, they're Indic, so there you have these strange ligatures tying these characters together. This is a great blog post. It kind of explains what's going on.

Steve: Yeah, the guy really dissected it, like pulled it apart for us.

Leo: Yeah, really good.

Steve: So, okay. This one, Leo, you're going to want to look on the next page here. The page starts with "EternalGlue." There's a picture of a network of computers. This is a little chilling. The NCC Group that we've talked about before, they're a well-known security group. They were contracted by an unknown or unnamed client.

Leo: I don't know, you can't really even read this. It's just it's a weird...

Steve: Right. But notice that it all starts in the upper left with a single machine. And what we're seeing is we're seeing an infection branching out within an organization. So as we know, EternalBlue uses the NSA's leaked SMBv1 exploit. They call this "EternalGlue," a rebuilt NotPetya. Remember that Petya and WannaCry were the two different pieces of malware based on EternalBlue.

So they write: "In June 2017" - so last summer - "we were asked by a client to rebuild NotPetya from scratch. Instead of the data destruction payload, they asked for telemetry and safeguards. Why? Because they wanted to measure what the impact of NotPetya" -

that is, essentially, a leveraging of the EternalBlue exploit - "would have been for them." They said: "We've completed the first phase of live testing in a secure environment deployed by our client." They said: "It has been a marathon, not a sprint. By the time we emerged from testing the code and the associated safeguards in December 2017, we had already been working with our customer in the lab for a number of months. This slow and steady approach has ensured everything works as intended, and the quality of telemetry is sufficient to answer the client's questions."

Then they wrote: "Christmas comes early: EternalGlue's first outing." And that's what this picture shows. They wrote: "On 7 December, EternalGlue got its first outing on the customer's engineering network," and so they said, i.e., a live network, but not corporate. Now, okay, if this is their engineering network, this is a big customer. So we don't know who this is. It's an unidentified client. But they said: "The result? More data than one could have imagined, and interesting insights as to the propagation in live environments."

So the headlines from Phase I of the experiment were, okay, now get this. The customer ran this EternalGlue, I don't want to call it malware, experiment ware on one machine in their engineering network with no privileges. It found three machines unpatched. It exploited those three machines to obtain kernel-level access. It infected those three machines. Within 10 minutes it had gone through the entire engineering network using recovered and stolen credentials. It then took the domain about two minutes later. One hundred and seven hosts were owned in roughly 45 minutes before the client initiated the kill and remove switch.

So here was a sandboxed, deliberately safe, carefully engineered leverage of, I mean, today the SMBv1, which has long since been patched by Microsoft, operating successfully within an existing organization. It found a couple machines which had escaped patching, got into them, used its position there to establish a beachhead, get into the kernel, get credentials, and then move through the network.

So if nothing else, this should be a chilling note that, if you have something as potent as the EternalBlue vulnerability is, you have to be really, really sure that you don't have a single machine that is vulnerable to it because it doesn't just let you get in. It lets you get down into that machine through that vulnerability, and from there maybe do much more damage, even in a network where other things are patched against it. So it just used EternalBlue, or in this case they called their system of course EternalGlue, in order to escape. But once it got kernel level on a couple of the machines it escaped to, in a short time it was in, infected 107 different machines in their engineering network. So interesting and certainly chilling. And I would say the client got their money's worth from the commissioning of that test by the NCC Group.

Just in passing, a fun story about a small bitcoin mining operation in someone's residence. T-Mobile complained that massive radiofrequency interference emanating from a local residence, and I don't know where it was located, at 700 MHz, was significantly interfering with the delivery of the company's cellular services.

Leo: You know, when they first came out with these high GHz processors, everybody was wondering. They're not broadcasting in broadcast ranges. Is this going to be a problem? It's the first time I've heard of it actually being a problem.

Steve: Yeah. The FCC's enforcement bureau said that they contacted the residence. I mean, I'm sure they drove a truck past and probably the antennas melted. They

contacted the people and said that continued use of - it was called the Antminer S4, and there was some confusion in the reporting because there's also reference to an Antminer S5, yup, there it is - that continued use of it will presumably...

Leo: It's hardware.

Steve: Yeah. Yeah.

Leo: Wow.

Steve: And what you do, of course, is you have...

Leo: A hundred of them.

Steve: Rack, rack, exactly, or thousands of them, rack after rack after rack.

Leo: This was in Brooklyn, so that's why the FCC - because the FCC has, like, five enforcement vehicles, so it's kind of amazing that they could find it.

Steve: Right. Well, of course T-Mobile was able to say our cell tower's gone down.

Leo: They knew where it was, yeah.

Steve: Yeah. They had a good idea of where it was. So they said continued use would be subject to fines, criminal prosecution, or seizure of equipment. So, yes, kids, you may, depending upon where you're located, may be disrupting cellular service within a perimeter of your mining operation.

Leo: So it's the processor? I mean, 700 MHz, is it the GPU? It's something - and it has got to be transmitting like crazy.

Steve: Well, yeah. Remember that any time, any time a wire has its charge rapidly changed, that will emit a bit of radio. So sine wave changes don't, but an edge produces a huge, high, large spectrum of interference. So basically, I mean, I know you know, Leo, that all of the computer cases that we've seen recently, they've got gasketing, and sometimes, well, even remember the Apple II, the inside of the case was like sprayed with that gray metal stuff. So you really have to go to some lengths to prevent radiofrequency from escaping from anything digital that is going fast. And of course mining equipment, the last thing they're thinking about is RF interference.

Leo: These are made in China. I doubt very much they have FCC "B" approval or

anything like that.

Steve: No chance. No chance. And now they're, what, at the T9? So that was the S4. Now they're up at...

Leo: Yeah, this is an old model.

Steve: Right.

Leo: You can buy them cheap on eBay. Pretty funny.

Steve: So a little bit under the cloud, or under the hood, rather, about Chrome's ad filtering. As our listeners may already know, what we talked about some months back that Google had announced Chrome was going to do was the so-called "intrusive adblocking." And I know you talked about this also over the weekend, Leo. It went live in Chrome on Thursday. And so from a technology standpoint I thought it'd be interesting to get some sense for what this is all about.

So Google explained the day before in a blog posting in their Chromium blog, like what was the basis for this, and what were they trying to do. It started with a survey of 40,000 Internet users throughout North America, which was taken by a group known as the Coalition for Better Ads, wanting basically to show them a bunch of things and rate how annoying different types of advertising was. The two most annoying were the so-called prestitial page, which covers where you are with a countdown before you're able to get through to it to the site. You know, I mean, Forbes does this; but I don't really mind because you're able to click past it, so it's like, okay. And then the second most annoying were the flashing animated ads.

So Google notes in their coverage of this, in their description of what they're doing in Chrome, that while some problematic ads are sourced by the advertising supplier, meaning, as we've talked about this often, a site creates a rectangle, like sets aside a rectangular area, and then the advertiser puts whatever they're going to put in there, meaning that the site doesn't control necessarily the content. They wrote: "The majority of problematic experiences, user experiences, are under the control of and at the specification of the site's owner," such as high advertising density and things like the prestitial page covers that is script running on the site that does this.

So Google writes: "This result led to the approach Chrome takes to protect users from many of the intrusive ad experiences identified by the Better Ads Standards [which is] evaluate how well sites comply with the Better Ads Standards, inform sites of any issues encountered, provide the opportunity for sites to address the identified issues, and remove ads from sites that continue to maintain a problematic ads experience."

I won't go through all the details that I have in the show notes. If anyone's interested they can look there. But essentially Google is giving sites, or has been giving sites, notice through the API if behavior that Google is seeing through Chrome is in violation and if Google, starting last Thursday, would be taking action against those sites based on this updated set of policies. And they conclude the posting saying early results are showing positive progress for users. Of course Google is couching all this as, look, we're not

wanting everyone to run adblocking. We're hoping we can take the pressure off of users taking their own actions by coming up with some compromise.

So they're saying - in their summary they said: "While the result of this action is that Chrome users will not see ads on sites that consistently violate the Better Ads Standards, our goal is not to filter any ads at all, but to improve the experience for all web users. As of February 12" - so that's, what, about two weeks ago - "42% of sites which were failing the Better Ads Standards have resolved their issues and are now passing. This is the outcome," they write, "we were hoping for, that sites would take steps to fix intrusive ad experiences themselves and benefit all web users."

They say: "However, if a site continues to maintain noncompliant ad experiences 30 days after being notified of violations, Chrome will begin to block ads on that site. We're encouraged," they conclude, "by early results showing industry shifts away from intrusive ad experiences, and look forward to continued collaboration with the industry toward a future where Chrome's ad filtering technology will not be needed."

So this strikes me as a good thing, but also it's a little scary. I mean, it's a web browser choosing to enforce a set of policies which it assumes are what its users want and changing the content of the sites. Now, it is the case that, when you go to a site where Chrome has made some changes, you will see a notification, typically a bar along the bottom, where you're able to opt out of Chrome's filtering of what it considers to be intrusive ads on that site. So the user has control. User has notification if Chrome has done this.

And it would be interesting to see how this goes. I mean, Google has been using, as we've often talked about on the podcast, their power, the power of being the majority browser on the web now, to make lots of changes in security. Now we're seeing some clear changes in advertising content, with appropriate controls and maybe with a good outcome, we can hope. Certainly the annoying ads are annoying to all of us.

Leo: Well, I notice Forbes doesn't put up that interstitial anymore. I don't know if that's just a coincidence. But I just went there.

Steve: Interesting.

Leo: Yeah.

Steve: Interesting, yeah. I just had one quick little tweet about my software, SpinRite, that's something I didn't know. Scott Napier - I saw this tweet as I was going back through my Twitter feed, pulling the show together. He said: "@SGgrc SpinRite has quite literally helped to keep the physical security of the Smithsonian intact, and I have convinced some skeptics of its power. Looking forward to 6.1."

Leo: Kind of a cryptic tweet there. What do you think?

Steve: So who knows? It may be that the security system for the Smithsonian uses hard drives that were having trouble, or the system went down, and it was like, ooh, crap, we need our security back up. Run SpinRite. And apparently in the process skeptics were

convinced. So thank you, Scott, for sharing that.

So I've decided, after reading an incredible number of suggestions for what we call the combination of Meltdown and Spectre, the single-word attempts, that there probably isn't a good one. But we do have one that is unique that hasn't been suggested yet.

Leo: Oh, boy.

Steve: So I'll call that the winner. Several people decided they liked S&M, which was what I finished with last week. But anyway, so just to remind people about how many different ways there are to put these words or word fragments together, since, I mean, just in going through the notes, my Twitter feed since last week. Nico de Smidt says: "How about Speckdown?" Life Cream Scoop: "How about Smeltdown?" Matt Fenton suggests: "What about Melspec, a twist on the word MIL-SPEC," for of course Meltdown and Spectre. Eat78, how about, he says: "How about SpecMelt?" Carol Saye suggests: "Meltspec."

But then I got, I finally saw one that was a little different that kind of hooked me. Again, I just think we're not going to end up with actually using it. But bcrypt.c - which is kind of an interesting handle, you know, bcrypt is a crypto. Anyway, he suggests "DownSpec'd," D-O-W-N-S-P-E-C-'-D, which is kind of cute because the consequence of this was that we had to reduce the specifications of our processors. They were downspec'd.

Leo: Oh, I like that, yeah.

Steve: So Meltdown and Spectre, DownSpec'd.

Leo: DownSpec'd.

Steve: So that's kind of the DownSpec'd attacks.

Leo: All right.

Steve: So anyway, of all of them, I think that one wins. I thank everybody for their suggestions. Please, no more. I think every possible combination of pieces of those words have been put together. But we did come across a new one. So thanks, everybody. Thank you, everybody. DownSpec'd, I think, of those I've seen.

Doug White, tweeting as @cpuguru, sent two things. He said: "Aloha, Steve," so maybe he's in Hawaii somewhere. "Just an FYI that I upgraded my Comcast Internet service to 1Gb speed last week. Trying to run the speed test resulted in about 460Mb throughput, even though the technician showed 1Gb at the port." He says: "I remembered an earlier podcast about the EdgeRouter X and that the internal throughput would be only about half a gig. So I purchased an EdgeRouter 4, replaced the EdgeRouter X. Now I have my 1Gb throughput." He said: "Figured I'd mention my experience as I imagine there are other home routers that will be likewise speed limited, even though they advertise 1Gb

ports."

Leo: Interesting.

Steve: And that's really - thank you, Doug. That's really good information because...

Leo: We had a radio - guy called the radio show, I think he was using Eeros, and he was getting 600Mb. So maybe that's it, yeah.

Steve: Yeah. So the fact that you've got 1Gb ports means that the wire can carry that traffic. But it doesn't mean that the complex routing and switching logic which figures out where to send each packet at what is really a very high rate, is able to keep up with the so-called line rate of the actual connections. And then in Doug's second tweet, which actually came in at a different time, but probably is a part of this work he was doing, he said: "If you're running a Ubiquiti EdgeRouter, looks like new firmware dropped last week."

So as a listener service announcement because I know that a lot of our listeners are using the Ubiquiti routers as a consequence of our affection for them, which is partly because they are so powerful from a configuration standpoint, you know, they're actual routers where you can give each of the ports of this little box its own network and subnetwork in order to create isolation. Just wanted to say that, I mean, and all that for, what is it, $49, that new firmware is available. So everybody may want to go check in with Ubiquiti and update themselves.

Two last tweets. Chip Steiner said: "When your ISP DNS has the best performance," and he says, parens, "(using DNSBench)," which of course is GRC's very popular utility for measuring DNS performance. He says: "When your ISP DNS has the best performance over Google DNS, Quad9, OpenDNS, but it uses DNS filtering, what's your advice?" And I don't really have any.

Leo: That's, well, I mean...

Steve: What, what?

Leo: You're going to opt for performance over security or injected ads; right?

Steve: Exactly. The DNS filtering is troublesome because it means that your ISP is in your business and, as you said, is able to perform various sorts of involvement. You would like when you enter a domain name that doesn't exist to get a DNS error. Instead, ISPs often use the opportunity to return a page with their own cert suggestions or their own advertisement of one form or another. So my sense is that using DNSBench is informative, I mean, it gave Chip a sense for the ranking of his options. But then try to see what it feels like in the real world because remember the DNS is cached. It's cached several levels within your own system. And the benchmark deals with both cached and uncached queries.

So I would say maybe try your ISP's DNS, get a feel for the way surfing the web feels, and then switch over, like to Quad9 or OpenDNS that is performing customer-facing filtering for security, which I think is very valuable, and see if you actually notice a difference. It's one thing for a benchmark to be able to find a difference. It's not at all clear what that means in terms of actual experience. So I guess I would say use the information from the benchmark, but then decide for yourself how you feel about the actual result either way.

Oh, and this is very cool. Chris Ryan sent me a tweet: "I saw this link for an IoT checker to see if your devices are on Shodan and thought I'd share." Okay, so here's what this does. This is very cool. You can query Shodan for your own IP to see if there are any entries in the Shodan database for your home network.

Leo: Oh.

Steve: Which would be good to know.

Leo: Yeah.

Steve: But this automates that: iotscanner.bullguard.com. So I-O-T-S-C-A-N-N-E-R dot B-U-L-L-G-U-A-R-D dot com. And there are two levels of tests. The first is a passive - so when you go to iotscanner.bullguard.com, it has your IP, in the same way that ShieldsUP!, GRC's test does. So it's then able to submit that IP using the Shodan API to see if Shodan reports any results. And so you're turning up clear on your network, which we would hope. That's good.

Then, if you want to, you can go to the next level and use this iotscanner.bullguard.com to actively scan your current public IP to see if it is able to find anything. And again you came up clean, Leo. And I did it both on my networks here and over at my residence, and they both came up clear at both levels of scan. So it's like, yes, good. But it's absolutely going to be the case that some of our listeners are going to discover something that doesn't make them happy when they do this. So iotscanner.bullguard.com. And thank you, Chris, for bringing that to my attention so I could tell everybody.

Okay. So Russian meddling technology.

Leo: Sounds like something you need [indiscernible] for.

Steve: Yes. So last Friday we know that Special Counsel Robert Mueller, his investigation into the question of Russian interference with the 2016 U.S. presidential election, released an indictment naming 33 Russian individuals and three Russian companies, accusing them of violating U.S. criminal law. I remember being sort of nonplussed by that. It's like, okay, well, so, okay. So the indictment charged the defendants with conspiracy to defraud the U.S., wire fraud, and identity theft.

And so, as we know, this is a technology, not a politics podcast. But what interests me, and I think is of interest to us here, is not and should not for this purpose the election politics aspect of the effort, but on the technical side of what was done to pull this off.

And I think this is relevant, not only for technical interest, but because we're hearing now a lot in the news and among the talking heads and commentators that Google, Facebook, Instagram, Twitter, and others were, or at least to some degree to be determined, were responsible, maybe negligent, that they should have known better than to have their services abused in this fashion.

So as I've looked at, you know, into what we learned, in many ways, as I mentioned at the top of the podcast, this struck me as being reminiscent of U.S. law enforcement and the U.S. FBI complaining about the use of encryption and the Internet going dark problem and demanding a golden key for access, my point being that it comes down to technology, and technology is a world that we understand. So the indictment alleges that Russians purchased servers located in the United States in order to obfuscate their origins, and then created and established hundreds of fake personas on social media which they carefully developed into leaders of public opinion.

So they established essentially IP addresses in the U.S. and then went about investing in the creation of social media personas. They used virtual private networks to open and operate social media accounts, and in that way behaving exactly as any U.S. citizen might. A lot of users use VPNs for various reasons in public settings, in hotels, in open WiFi settings, in order to get security and privacy. And they also found and allege that Russian agents stole U.S. identities to open accounts with PayPal to further substantiate and support these false personas and identities and used those to purchase advertisements on social media sites.

So the bottom line is, leveraging the technology of the Internet, which is about freedom and openness and anonymity to varying degrees, they were able to convincingly pose as U.S. citizens and also, again, just by nature of the way social media works today, they were able to recruit through this technology real paid Americans to also engage in political activities, to promote political campaigns, stage political rallies. And those Americans that had been recruited had no idea nor any reason to suspect that they were communicating with Russian agents. I loved your accent verbally at the beginning of this, Leo, because it occurred to me that there's no accent in a tweet.

Leo: No.

Steve: There's no way to know how good, I mean, as long as the written English is convincing, there's an inherent disconnection. So interestingly, the indictment alleges that in some cases the Russians fomented unrest on both sides of the political divide at the same time. It said that, quote: "After the election, the defendants allegedly staged rallies to support the President while simultaneously staging rallies to protest his election. In one instance the Russian defendants organized one rally to support the President-elect and another rally to oppose him, both in New York on the same day." So there was just a concerted effort to stir things up.

So here we are. We in the U.S. invented the Internet. It was our technology. And we're collectively profiting mightily from its existence, largely because it's such a wide-open communications platform. But what happened was that an adversarial country very cleverly used our own technologies and even our own bandwidth, on our soil, against our nation's interest.

Now, what I liked in this was something that Mike Masnick at Techdirt wrote that I think helps to nicely frame sort of like where we go from here. He did a posting titled "DOJ Russia Indictment Again Highlights Why Internet Companies Can't" - can't, C-A-N-'-T,

cannot - "Just Wave a Magic Wand to Make Bad Stuff Go Away." He says: "This was not just some run-of-the-mill 'pretend to be Americans.' This was a hugely involved process to make it very difficult to determine that they were not Americans."

He writes: "I've seen some people online claiming that this shows why the platforms have to take more responsibility for who is using their platform." And then he quotes a tweet from a Renee DiResta, who tweets from @noUpside. On February 16th Renee tweeted: "While you read the Mueller #Indictment, remember the tech CEO mantra: 'We don't want to be the arbiters of truth.'" The tweet goes on: "These platforms were used exactly as they were designed to be used. Here we are a year later, and still no accountability or governance," ends this tweet.

So Mike continues: "But my read on it is exactly the opposite. It shows just how ridiculous such a demand is. Would any of us," he asks, "be using these various services if we were all forced to go through a detailed background check just to use a social media platform? That seems," he writes, "excessive and silly. Part of the reason why these platforms are so useful and powerful in the first place is that they're available for nearly everyone to use with few hurdles in the way. That obviously has negative consequences in the form of trolling and scams and malicious behavior, but there's also a ton of really good stuff that has come out of it."

He says: "We should be pretty cautious before we throw away all the value of these platforms just because some people use them for nefarious purposes. People are always going to be able to hide their true intentions from the various platforms, and the response to that shouldn't be put more blame on the platforms, it should be a recognition of why it's so silly to blame the tools and services for the actions of the users."

He finishes: "Yes, we should be concerned about foreign attempts to influence our elections, while noting that the U.S. itself has a long history of doing the same damn thing," he writes, "in other countries, so this is a bit of blowback. But blaming the technology platforms the Russians used seems to be totally missing the point of what happened and risks making the Internet much worse for everyone else."

And I can't find much argument with that position. Because when, I mean, maybe you can apply sophisticated AI, heuristics, or maybe do a somewhat better job. But it's not at all clear how somebody who has an agenda, even if it's foreign-inspired, differs from an American who has an opinion.

So anyway, I thought it was interesting that they're using technology, given that this indictment and these allegations are true, using technology in a way that it was designed, essentially residing in the U.S. and acting for their own ends. And it's not at all clear how it's easy to do more than simply jump up and down and stamp our feet and say we want it to be different. Yeah. But how do you pull that off? So there.

Leo: You're not drawing me into this one.

Steve: No? Okay.

Leo: I can see you waiting.

Steve: And that's our podcast, Episode 651.

Leo: I can see you waiting.

Steve: For February 20th, 2018.

Leo: It's a very difficult problem.

Steve: It is a problem, and it's one we've got. It'll be interesting to see what happens moving forward.

Leo: You don't want to limit free speech.

Steve: No.

Leo: I'd be a little more sanguine about it if Twitter and Facebook had been a little more forthcoming. They haven't been helpful at all.

Steve: Yeah.

Leo: But you're right, I mean, I don't know exactly how, you know, you don't want a proof of identity. I don't know. I don't know. And obviously the proof of identity, they got around it with identity thieves.

Steve: Yes, exactly.

Leo: Well, Steve, we do this show, and there's plenty - and I'm sure you'll hear from people on both sides of the equation. But there's plenty more to talk about. Every Tuesday, 'round about 1:30 Pacific, 4:30 Eastern, 21:30 UTC, we convene and discuss the latest in security news with this guy right here, Steve Gibson. You'll find him at GRC.com, the Gibson Research Corporation. While you're there, pick up his bread and butter, SpinRite, the world's best hard drive recovery and maintenance utility. If it's good enough for the Smithsonian, it's good enough for you: GRC.com.

Steve: And you know that it was on the International Space Station, too.

Leo: SpinRite was?

Steve: Yeah, a copy was sent up some time ago. They liked it because they were literally counting bytes, and it was so small as a recovery utility that they were able to stick it on like a floppy with a whole bunch of other stuff, like utilities that they needed, rather than it needing its own megabytes of space. And so, yeah, it was used on the International Space Station for quite some time.

Leo: Very interesting. I did not know that.

Steve: Speaking of famous locations where SpinRite has been used.

Leo: Well, you ought to have it in your house. If you've got hard drives, you need SpinRite. Also he has the show there, of course, 64Kb MP3s and transcripts. Elaine Farris writes those all out in whatever order Steve refers to things, even when he messes it around. You could just find that, read along as you listen, at GRC.com.

We have video as well as audio at our site, TWiT.tv/sn, for Security Now!. And of course you can always subscribe on your favorite podcast appliance, and that way you'll get every episode the minute it comes out. Collect all 651. You need the complete set. Hey, it's just bits. Come on. There's plenty. You're not in the space station, for crying out loud.

We will be back next Tuesday. But in the meantime I think you have a little more time to take our TWiT Survey, if you want to. We don't collect information about our users in any other fashion, but once a year we like to do a little short survey and find out whatever you're willing to tell us about yourself, like you like big giant cups of coffee, perhaps. I don't know if we ask about that. What coffee cup size would you prefer? Maybe we should add that to the next quiz, next survey.

Steve: That works for me. How many grams of caffeine do you consume a day?

Leo: TWiT.tv/survey. Don't forget, it's very easy to listen to our shows on your voice-activated device, no matter who makes it - Apple, Google, even Microsoft, and of course Amazon. Just ask for Security Now!. Say, hey, you know who, "Hey, Echo, I want to listen to Security Now!," and it will play you the most recent episode. You can often listen to the live stream as well. Just say, "I want to listen to TWiT Live."

Oh, and we're also on the Flash Briefing. Many of our shows we just do little one- or two-minute clips that, if you listen to the Flash Briefing on your Amazon Echo, go to the Echo app on your phone and add TWiT to your Flash Briefing rundown. I think I've said enough. I'm going to shut up now and wish you all a wonderful evening. And we'll see you next time on Security Now!.

Steve: Thanks, buddy. Bye.