SECURITY NOW!

Transcript of Episode #654

AMD Chipset Disaster

Description: This week we discuss the just-released news of major trouble for AMD's chipset security, ISPs actively spreading state-sponsored malware, Windows 10 S coming soon, a large pile of cryptocurrency mining-driven shenanigans, tomorrow's Pwn2Own competition start, surprising stats about Spam botnet penetration, and a Week 2 update on the new Memcached DrDoS attacks.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-654.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-654-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson's here. He was all set to talk about the Memcached exploit spreading like wildfire. That was the culprit in those DDoS attacks, huge DDoS attacks last week. But then, all of a sudden, some significant flaws were revealed in the AMD chipsets, and Steve decided he was going to do a deep dive on that. Lots more security news, too. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 654, recorded Tuesday, March 13th, 2018: AMD Chipset Disaster.

It's time for Security Now!, the show where we protect you and your loved ones online. This is becoming one of our biggest and best shows, and it's thanks to this guy right here, Steve Gibson of GRC.com. Hello, Steve.

Steve Gibson: Hello, Leo. Great to be with you again, as always, for our 654th episode of Security Now!. I had intended for this to be a follow-up on last week's disclosure and discussion of the big memcached huge, record-breaking, 1.7TB distributed reflection denial of service attacks. And then something even worse happened. So I was going to call this "Memcrashed Update" or something, and I have all of the sort of like what's happened since last week, which is a lot on that front. And then this morning a security research firm in Tel Aviv who gave AMD 24, count them, 24 hours of notice...

Leo: Oh, boy.

Steve: Yeah, dropped a major bomb that they had found 13 remotely exploitable vulnerabilities in the chipset that AMD uses to sort of offload its security stuff, sort of like the Secure Enclave, and also just a chipset that's like a glue chipset. Anyway, this, I mean, the podcast was pretty much assembled, and everything was in place. And I thought, okay, good. And then, because the press release hit and the Internet went nuts, first I was thinking, okay, well, there's not time to know about all this yet, so let's just do that next week.

But as I kept reading it, like getting into more of this, and the whole research paper is online, so there was nothing I couldn't find out. And it's like, oh. We can't wait. So this suddenly got renamed from "Memcrashed Update" to the "AMD Chipset Disaster," which surprisingly is not hyperbole, yeah.

Leo: Doesn't sound good. Oh, man.

Steve: So we're going to discuss the just-released news of major trouble for AMD's chipset security. And the only possible mitigating factor I can imagine is that maybe there aren't many of them out there. Except I think there are because they're embedded, and they're laptop, and they're server, and they're enterprise. And it's like, ouch. But we're also going to talk about ISPs in East Asia have been found actively spreading state-sponsored malware, that is, ISPs themselves.

A note about Windows 10 S that I noted you and Mary Jo and Paul were talking about last week. I just wanted to let our listeners know about Microsoft's announcement about 10 S. We've got a large pile of cryptocurrency mining-driven shenanigans because of course there's money in them thar hills, and so you're going to have some shenanigans. Also tomorrow is our annual Pwn2Own competition start. But there's some news about that, and a bit of a twist, about a set of contestants that will no longer be appearing. I have some surprising stats, I mean, stunning stats about where our spam is coming from.

And then we will finally get to what was going to be the topic of this week, but now it's got pushed down the stack, and that's an update on the memcached distributed reflection denial of service attacks and what's going on with them. And in fact that's our Picture of the Week that we'll talk about in a second. So I think, yes, another interesting and, for owners of AMD-based systems, maybe a little frightening, too.

Leo: Is it the newer Ryzen systems, or is it older AMD?

Steve: Yes, the newer Ryzen...

Leo: It's Ryzen systems. Oy.

Steve: The Ryzen and the EPYCs, I think, is that the way you pronounce it?

Leo: Yeah. I think a lot of people bought these Ryzen systems.

Steve: Yeah.

Leo: That's not good.

Steve: Yeah.

Leo: Well, stay tuned. The thrilling, gripping details of the latest bad news exploit. Steve?

Steve: So our Picture of the Week was going to tie into the title of the podcast.

Leo: Didn't have time for a second illustration, huh?

Steve: Which changed at the last minute. Okay, but get this. More than 15,000 memcached DDoS attacks were aimed at 7,100 sites in just the last 10 days.

Leo: Wow. This can't all be the same person; right?

Steve: No. And this is the problem is that there are, we're not exactly sure how many, but at least 17,000 open reflectable memcached servers, which anyone on the Internet can use. And as we'll be talking about, if we survive till the end of the podcast, there are now easy-to-use tools where you enter an IP address, and it goes and finds available memcached servers from Shodan and starts filling them up and sending them reflection UDP packets. I mean, it was no exaggeration a week ago when we made the podcast about this because this is without question the most potent class of DDoS attacks we've yet encountered. It is due to the fact that a tiny query can produce a massive response which can be spoofed. The amplification factor, and we've talked about amplification factors of, oh my god, 200. And so you send in 10 bytes, and you get 2,000. Whoo. Okay, this amplification factor on these memcached DDoS attacks is 51,000.

Leo: So for every byte you get 51K.

Steve: Yes.

Leo: Wow.

Steve: Yes. So anyway, we'll come back to this at the end of the podcast. But what this is demonstrating is that this is all, I mean, if anything could have pulled hackers' attention away from cryptojacking, it would be DDoS attack capability that is this juicy and potent and capable. So it's like, great. So I'm sure we'll be talking about it next week. The industry is trying to rally itself; but, as we know, there are servers scattered all over the Internet. Many of them are unattended. Many of them are in closets and have been forgotten, and no one is paying any attention to them. And they'll just happily accept huge blobs of data and then send them out on request. So, yeah. Anyway, there's lots more news that we will get to at the end of the podcast. But what preempted that as our main focus...

Leo: Must be pretty bad if it preempted that.

Steve: So a brand new, some of the first, the very first coverage confused this with Spectre and Meltdown, and this has nothing to do with Spectre and Meltdown. We know that AMD chips have never been susceptible to the Meltdown problem, but that they do need some microcode fixes in order to help them mitigate the Spectre problem. So there's that. But a brand new, completely different, and just coincidental timing, set of 13 critical flaws have just been reported to affect AMD's Ryzen and EPYC processors. You may want to pull up this site, AMDflaws.com. I mean, these people are not being shy about their reporting. They've got...

Leo: They registered a domain name. They're serious. That's serious.

Steve: They registered a domain. They've got graphic artists at work. They've got videos. They've got...

Leo: But this is not AMD. This is the group that discovered it.

Steve: Correct.

Leo: Is that right?

Steve: The group is, yes, the group is CTS Labs. Oh, and you can imagine AMD is not happy right now. They were given, as I mentioned at the top of the show, one day, 24 hours' notice of this all going public.

Leo: But this has a good name, which means it's going to spread like wildfire. Because it's all about the name; right? Ryzenfall? I like it.

Steve: That's actually a brilliant name. So the group is CTS Labs. They're an Israeli Tel Aviv-based team of researchers. Their press release dated March 13th, which is today, 10:00 a.m. Eastern, said: "CTS Labs, a cybersecurity research firm and consultancy, today released a severe security advisory on Advanced Micro Devices, Inc. processors. A CTS Labs security audit," they write, "revealed multiple critical security vulnerabilities and manufacturer backdoors in AMD's latest EPYC [E-P-Y-C], Ryzen, Ryzen Pro, and Ryzen Mobile processors." Now, technically these are in the chipsets, that is, the supporting chips. But we'll get to that in a second.

"These vulnerabilities have the potential to put organizations at significantly increased risk of cyberattacks. CTS Labs has produced a whitepaper report further detailing these vulnerabilities available at AMDflaws.com." And there's, like, several offers on that page to download the whitepaper, which of course I did, and spent some time with it.

"CTS Labs has also shared this information with" - I'll add editorially yesterday - "AMD, Microsoft, HP, Dell, and select security companies, in order that they may work on developing mitigations and patches, and examine and research these and other potential vulnerabilities at the company," that is to say, AMD. "CTS Labs has also shared this information with relevant U.S. regulators. CTS Labs," they write, "is a cyber-security research firm and consultancy based in Tel Aviv, Israel specializing in hardware and embedded systems security. For more information about CTS Labs, please see cts-labs.com."

Okay. So as I've said, the vulnerabilities don't affect AMD's Zen CPU cores themselves, but rather two different chips that are part of the Ryzen and EPYC system set. The first is the ARM-based AMD Secure Processor, which as I mentioned before is kind of like their equivalent of the Secure Enclave; and the second is known as the Promontory chipset, which is produced by Taiwanese subcontractor ASMedia. And these people are not happy about the work that ASMedia, that is, CTS Labs is unhappy about that.

So all there has been so far from AMD because, I mean, this just happened, is that AMD responded: "At AMD, security is a top priority, and we are continually working to ensure the safety of our users as new risks arise." Daily. They continue: "We are investigating this report, which we just received, to understand the methodology and merit of the findings." And so I need to say also that, I mean, my first reaction upon seeing this was, okay, wait, maybe we should hold off saying anything. But that would mean a week delay rather than a day delay. And so I thought, well, no, the best that could happen is that it turns out that this is wrong. But the more I looked at it and dug into their research, the more credible and verifiable this appeared.

Okay, so what happened? Thirteen critical security vulnerabilities and what they describe as "manufacturer backdoors" were discovered throughout AMD Ryzen and EPYC product lines. In their little Q&A that is at the AMDflaws.com site they said: "Am I affected?" And the answer: "Any consumer or organization purchasing AMD servers, workstations, or laptops are affected by these vulnerabilities." And of course we have to back that off a little bit and say, if they contain these affected chipsets. And so they said: "This site is to inform the public about the vulnerabilities and call upon AMD and the security community to fix the vulnerable products." I doubt that much calling upon will be necessary. But they have another site, safefirmware.com, where the whitepaper can be found at /amdflaws_whitepaper.pdf.

So there are a total of four vulnerabilities, three in firmware and one in hardware. And this Chimera is the fourth hardware vulnerability which is unfixable, they claim, because it is a hardware problem. It is something, apparently a backdoor built into, presumably by this AMD subcontractor we'll get to in a second. The other three - there's Ryzenfall, there's Fallout, and MasterKey - are firmware vulnerabilities which presumably can be fixed. These guys in their research did succeed in rewriting the firmware remotely of this firmware reprogrammable outboard chip that does the security and were able to install persistent threats into the firmware.

So it looks to me like the biggest problem, I mean, just in terms of practical remediation, is with Chimera. And so they write: "Backdoors Inside Ryzen Chipset: The Chimera vulnerabilities are an array of hidden manufacturer backdoors inside AMD's Promontory chipsets. These chipsets are an integral part of all Ryzen and Ryzen Pro workstations. There exist two sets of backdoors, differentiated by their implementation. One is implemented within the firmware running on the chip, while the other is inside the chip's ASIC hardware. Because the latter has been manufactured into the chip, a direct fix may not be possible, and the solution may involve either a workaround or a recall.

"The backdoors outlined in this section" - I'm reading from their paper - "provide multiple pathways for malicious code execution inside the chipset's internal processor. Because the chipset is a core system component, running malware inside the chip could have far-reaching security implications. The diagram below" - and they have a diagram in their PDF that basically shows the Ryzen CPU connected to this big red box, red because bad, where, like, everything, all of the I/O in the system is going there, the LAN and the WiFi and the USB, I mean, it's the switchboard for the CPU.

So it says: "The diagram below was taken from the instruction manual of ASUS Crosshair VI Hero Ryzen motherboard. It can be seen that not only is the chipset connected to the computer's USB, SATA, and PCIE ports" - which is to say pretty much everything - "it is also linked to the computer's LAN, WiFi, and Bluetooth. In our research," they write, "we have been able to execute our own code inside the chipset and then leverage the latter's Direct Memory Access (DMA) engine to manipulate the operating system running on the main processor. These two capabilities form the foundation for malware and provide a proof of concept. We believe that, with additional research, a determined attacker" - who of course now knows where to look - "may also be able to reach the following capabilities."

And so they tick off a key logger. It may be possible to implement a stealthy key logger by listening to USB traffic that flows through the chipset. Network access, it may be possible to implement network-based malware by leveraging the chipset's position as a middleman for the machine's LAN, WiFi, and Bluetooth. Bypass memory protection, it may be possible to leverage the chipset's position to access protected memory areas such as the system management RAM. They say: "We have verified this works on a small set of desktop motherboards."

So then they have a second section that I'm sharing just because it's inflammatory but worrisome. "Third-Party Chip Design Plagued with Hidden Backdoors" is the topic for this. They said: "In November 2014 it was announced that AMD signed a contract with the Taiwanese chip manufacturer ASMedia, according to which ASMedia would design AMD's chipset for the upcoming Zen processor series. This chipset, codenamed Promontory, plays a central role within the company's latest generation Ryzen and Ryzen Pro workstations. It is responsible for linking the processor to external devices such as hard drives, USB devices, PCIE cards, and occasionally also network, WiFi, and Bluetooth. Although it is branded AMD, the Promontory chipset is not based on AMD technology. Rather, it is an amalgamation of several integrated circuits that ASMedia has been selling to OEMs for years, all merged together onto a single silicon die."

Okay. I'm skipping a whole bunch where they really tear ASMedia apart, saying that the chips that were glued together have a long history of problems. They exist on other motherboards and may themselves be problems. But essentially the essence is that ASMedia just pulled a bunch of their existing stuff, kind of threw it together, stuck it on a die, and said here you go.

And so in their Q&A they ask: "Doesn't this publication put users at risk?" And I would argue their statement. "No," they say, "all technical details that could be used to reproduce the vulnerabilities have been redacted from this publication." They say: "CTS has shared this information with AMD, Microsoft, and a small number of companies that could produce patches and mitigations." On the other hand, we know from a practical standpoint that telling the world 24 hours after you reveal this to AMD, there's just no way that that's not irresponsible. At a minimum, 90 days to allow AMD - or a week to allow AMD to evaluate it independently, come to their own conclusions, and produce a response. Not even that courtesy was given.

So it's very difficult not to hold these people responsible for, to some degree, any damage which occurs as a consequence of the fact that it's not possible to expect anyone to react in 24 hours to this kind of disclosure. And this is an outboard chipset. We know very little about it at this point. I'm sure I'll have a follow-up next week. The good news is they did not release, I mean, thank goodness, proof of concept or additional details.

The problem is that the world now knows, the hacker world now knows there is a big juicy target here. And there's now a race to see, well, first of all, whether there's any remediation that can be done. Then that has to get created and has to get somehow pushed out to the affected motherboards while, in parallel, bad guys are working to figure out what it was that CTS has figured out and leverage it before that can happen.

So anyway, there's no takeaway for our listeners. Nothing we can do at this point. I have a bunch of AMD-based machines. I'm largely an Intel shop, but I have AMD because they're a strong source of 64-bit systems, as we know. So we'll just be following this news and of course keep our listeners informed. The good news is it's not clear what role Microsoft could play. We do know that Microsoft has an active patchable channel, a patching channel. As we talked last week, they have started to use that to patch the Intel microcode.

This is outboard microcode, so this will be microcode of that chipset. So it's not clear to me what the path is for getting that fixed. Maybe, if there's access to it from the main CPU, that is, programmatic access, then presumably AMD could produce and Microsoft could push a patch for Windows users, at least. And of course Windows is not the only OS to run on AMD chips. Maybe we'll see something happen in the Linux community even before. I mean, who knows?

We'll be watching it closely because this is - as you commented, Leo, at the top of the show when you heard it was the Ryzen family that were affected - this has got to be really bad news for AMD. I'm just - I'm disappointed, frankly, that CTS did this. It would have been much more reasonable for them to give AMD some time to respond, for example, for a remediation and a patch to be available in the channel, and then for these guys to say, okay, we're the people who found this. I don't understand why that didn't happen.

Leo: Unless they felt like it was going to imminently come out in other ways or...

Steve: Yeah, good point, yup.

Leo: And there are reasons you might want to make it go public as soon as possible.

Steve: Well, I mean, nothing could possibly put AMD under more pressure. So if there was any sense that they weren't going to be listened to, that AMD was blowing them off, that it wasn't going to get fixed, like as you said, maybe they do have some information about something else going on that really makes this have to happen as fast as possible. We just don't know yet. So I imagine more information will be coming. And we'll, of course, keep our listeners in the loop.

Leo: Yeah, wow.

Steve: Yeah. Not what you want to have happen if you're...

Leo: I guess this means you shouldn't buy an AMD processor for a while. I mean, there are people in the chatroom saying, "Gosh, I was going to go get one today."

Steve: Yeah.

Leo: I would wait a little while, at least to hear what mitigations there are. Maybe it's not a real threat. Maybe AMD will come out and say, oh, we've analyzed this, it's not a real threat. Or I don't know.

Steve: Yeah.

Leo: It's not like Intel doesn't have its own problems; right?

Steve: Right. And maybe, I mean, most of this is firmware updatable, apparently, except for the one, the hardware variant of the Chimera flaw. And we don't know that you can't work around that with firmware to maybe shut down something. Maybe it means you have to turn off one of your USB ports, or who knows. But it would be difficult, yes, Leo, to go and lay money out right now. Well, but it's also a function of your exposure. If it's your own machine...

Leo: Do we know what the mechanism for an attack would be?

Steve: No, we don't yet. Although, based on what little we know, it looks like because this chipset, like that outboard chip, this Promontory chip, which is where the problems are, it is the glue for all the I/O. So hard drive, LAN, WiFi, Bluetooth, I mean, you name it, it potentially has access into the chip. Though, again, we don't know how it is that you leverage this. And the good news is, I mean, AMD just must be gasping right now. I'm sure within the next day or two, because I'm sure they're motivated to help create some AMD-side calibration, we should know more.

And so maybe wait just a couple days and see what AMD says. If they're able to say, okay, definitively, this is only a local attack through USB, that kind of thing, then you would probably - and if they're able to say "and we've got a fix on the way," then I'd say, yeah, fine, you know, go ahead. Wow.

Leo: Poor Ant Pruitt just got his a week ago. But it doesn't necessarily mean this is panic time for you, either, as a user.

Steve: No, well, and here we are, what is it, it's the middle of March, and we've been talking about Spectre and Meltdown all year. And there still isn't a widespread known exploit of that. And so, yeah, it's not, you know, we're 10 weeks in. And the good news, I mean, if there's an upside, it's that this couldn't possibly motivate AMD any more strongly. And no doubt ASMedia is saying whoops.

Leo: Yeah, no kidding. Well, I even wonder, whoops or did they get caught up to no good? Right?

Steve: Yeah.

Leo: If you put a backdoor into these chips...

Steve: If it's really a hardware backdoor, it's very hard to, I mean, that AMD was unaware of, that AMD didn't ask for...

Leo: Right. This may not be AMD's fault. Probably those, at least, are not AMD's fault. Probably ASM just did it. Is ASM on mainland China?

Steve: They're Taiwanese.

Leo: Okay. So not Taiwan. That's different.

Steve: So, yeah, not mainland. But who knows?

Leo: I would be suspicion of mainland Chinese manufacturers because they could easily be subverted by the government, whether they wanted to or not.

Steve: Well, in fact we've got a couple stories about what's happening with China today. So we'll be talking about that in a second. So, yes, you're certainly right.

One thing that is a little disturbing - this doesn't affect people outside of Turkey, Syria, and Egypt. But from a technology capability standpoint it is worrisome because it could affect us all. And that is that ISPs in those three countries have been caught using - I'm annoyed with this term "deep packet inspection." All it really means is hardware which is operating at the packet level. That is, it's not just a router which is looking at packet headers and sending things around. It's going into the packet in order to see what's going on. It turns out that there is a company, Sandvine, which has a line of hardware called PacketLogic devices, which Sandvine argues are not designed for this purpose, and they're not happy that they've been called out and having been found.

Leo: Well, we should point out that there are ISPs who use Sandvine for deep packet inspection.

Steve: Yes.

Leo: And I think it was Comcast, it was one of the U.S. ISPs was using it to block

BitTorrent.

Steve: Yes, yes, exactly. And in fact Citizen Lab are the researchers who caught this behavior and reported their findings to Sandvine, which says that Citizen Lab's, who I strongly believe is correct, report is false, misleading, and wrong, and demanded that Citizen Lab return to them the secondhand PacketLogic device which Citizen Lab used to confirm the attribution of the network fingerprint that the Sandvine devices were using. So it looks to me like Citizen Lab did everything right.

So here's what they found: Turkey's Telecom network, which had some overlap in Syria, so thus some Syrian ISPs were involved, or Syrian citizens, at least, were using Sandvine's PacketLogic devices to redirect hundreds of targeted users - so not everybody, but if you were a known customer of Turkey Telecom, so for example journalists, lawyers, and human rights defenders, your typical targets in those environments - to malicious versions of legitimate programs, and those programs were bundled with FinFisher and another spyware or malware known as StrongPity, when they tried to download them from official sites.

Now, that's what's interesting is unfortunately there are a number of sites, many of which are HTTPS. But when you go to download their software, the download link is HTTP. So what happens is you're at a legitimate site - and we're talking, for example, Avast Antivirus; CCleaner unfortunately got caught up in this again, and remember that they were a victim of some compromise not too many months back; Opera; and 7-Zip. Those sites, at least at the time of this reporting, were downloading their software over non-HTTPS. So even if you went to those sites over HTTPS, so you knew you were there, the download link was not.

So what this PacketLogic device allowed was that to be intercepted since, without TLS, without being in a TLS tunnel, you have no authentication of the connection. And that caused the legitimate, the otherwise legitimate links and their programs to be funneled from another source undetected, and people believed they had done the right thing. They believed they were, you know - they were at the legitimate site, but the actual software they received had been tampered with. So that's the case in Turkey.

In Egypt, the same Sandvine PacketLogic devices were being used by an Egyptian ISP to make money in two different ways. One, they were secretly - the ISP was injecting cryptocurrency mining script into every HTTP web page their own users, their own customers visited, in order to mine Monero cryptocurrency. So again, non-HTTPS. The Sandvine PacketLogic devices were unable to deal with that due to its encryption. But the second you dropped out of HTTPS, as soon as you just started using a plaintext connection, then they were able to manipulate your traffic and, in this case, add some cryptocurrency mining JavaScript into any non-HTTPS websites that their client customers, that their customers visited. And they were also redirecting some users to web pages with affiliate ads as another way of generating additional revenue.

So anyway, as far as we know, that's not happening to users outside of those countries. But it is a real heads-up that there's a vulnerability which is exploitable and we know has been exploited - if not against us, at least theoretically, well, it has been exploited against those customers, those targeted customers, leveraging the fact that, without TLS encryption on the connection, you are still subject to manipulation, even if you're at an otherwise HTTPS site. And, boy, we have to, you know, maybe it's because they're using content delivery networks; they're using a CDN that doesn't offer HTTPS services. I've not looked any further to see, for example, why legitimate companies like Avast, CCleaner, Opera, and 7-Zip might be offering non-HTTPS content in some cases. It's worth keeping an eye on.

Leo: All right, Steve. Continue, my friend.

Steve: So we've not talked about Windows 10 S because we've not talked a lot about...

Leo: We know what you think of Windows 10. We know what you think.

Steve: Yes, we do. So I did want to just mention, so that it hadn't been ignored on this podcast, that Microsoft has announced a change in their policy relative to Windows 10 S.

Leo: Paul would say this was what it always was, but they just didn't communicate it well.

Steve: Okay. And now they are. Last Wednesday evening - is it the Corporate Vice President in charge of Windows, is it...

Leo: Joe Belfiore, yeah.

Steve: Belfiore, okay, yeah, Joe Belfiore.

Leo: We call him Joey B. He's great, actually.

Steve: Joey B. He wrote a short blog post to describe what's happening that I'll just share, just so that we have it on the record here. So it was Windows 10 in S mode coming soon to all editions of Windows 10. And that's why I think it's relevant to our listeners is that it's becoming a "mode" for Windows 10.

Leo: Right.

Steve: So he said: "Some of you may have seen a discussion around our plans for Windows 10 S on Twitter today," speaking of last Wednesday. "And given some additional questions I've received," he writes, "I thought it might be helpful to share more about our plans with Windows 10 S. Last year we introduced Windows 10 S, an effort to provide a Windows experience that delivers predictable performance and quality through Microsoft-verified apps via the Microsoft Store. This configuration was offered initially as part of the Surface laptop and has been adopted by our customers and partners for its performance and reliability."

Leo: Yeah. I had that Windows Surface laptop. And what they offered at the time was great. You could just push a button and turn it into Windows 10 Pro.

Steve: Right, meaning you could un-S yourself.

Leo: Un-S yourself, yeah. Kiss your S goodbye.

Steve: Yeah, that's right. And apparently everybody did.

Leo: Yeah.

Steve: He said: "Since that time" - well, and I've heard Paul say it was basically unusable.

Leo: I would disagree. But if you can't put Chrome on them - the things about S were it had to be, I think it had to be you could only use Store apps.

Steve: Correct, for sure.

Leo: I think they had to be 32-bit. I can't remember that issue. But the problem is...

Steve: Probably have to be the UWP...

Leo: UWP, yeah. So you couldn't use Chrome. You couldn't use LastPass. I mean, there was a lot things - I guess actually you could use LastPass, but you couldn't use Chrome. You couldn't use...

Steve: Well, anything else that you wanted to use that...

Leo: ...anything that wasn't in the store; right. So the real question, which we asked on the show, is...

Steve: So sort of like the Windows Phone of Windows.

Leo: Well, it wasn't that bad. You get most of the stuff. It's just the real thing is, is there really any benefit to it? Is it more secure? Is it faster? Does it have better battery life? That was never clear. That's what Microsoft asserted.

Steve: And what do you think their intent was? Because I've never really understood. I mean, did they want to create a monitored, managed enclave like for...

Leo: No, it wasn't that sophisticated. That's the problem. It was basically Windows 10 where you couldn't get anything but what was in the Store. Kind of like if you took your Mac and you said only allow apps from the App Store, which is [crosstalk] OS X.

Steve: That's probably what "S" stood for was Windows 10 Store Edition; right?

Leo: Yeah. And I think it was to compete with the Chromebook. Even though the Surface laptop was not inexpensive, I think it was assumed that all the other computers that ran Windows 10 S would be for education, where you might want to lock it down, and for less expensive environments. It was to compete with the Chromebook. But the truth is, yeah, it was kind of a muddy product to begin with. And so considering it a service or a mode of Windows 10 makes more sense. If you had a switch, just like you do in OS X, where you say we'll only allow people to buy stuff from the App Store, don't allow them to download, it would be more secure; right? If you couldn't download stuff that wasn't in the App Store?

Steve: Absolutely. Absolutely.

Leo: So that was the idea.

Steve: Well, for example, I'm sure that everything in the App Store has been digitally signed.

Leo: Right.

Steve: And when I was talking just a minute ago about ISPs intercepting HTTP and doing packet-level redirection in order to cause malicious code to get downloaded, I meant to say that just two days ago I was looking for, because it's me, an old version of Adobe Reader. I don't want DC or 11 or any of the new nonsense where you have to sign up and subscribe and join and be a member. I just want a PDF to open. And so I ended up finding, I think it was a version 11, or maybe it was 9.

Anyway, I found something that looked right, but it wasn't from Adobe. Since then I've realized that Adobe's FTP server does have everything that they ever did there. It's like, oh, thank goodness. But a couple days ago I found one from what looked like a reputable download site. And I'm not recommending that anybody do that because I was lucky, I felt I was lucky to find it. What I downloaded was 35MB, and I was very careful not to run it. Instead I checked to see if it was signed, and it was signed by Adobe. I verified the certificate. I verified the signature. I mean, I did all this other stuff that I know how to do, and I know that our listeners know how to do. Most people don't.

But the point is that that's what you have to go through because in our current world there's just so much crap on the Internet that is trying to take advantage of you. So it's certainly an advantage for a neophyte Windows user to, if you turn on S mode, then Windows is certainly going to not let you download anything from other than the Store, and anything from the Store will be digitally signed by who knows.

Leo: And at least presumably vetted, although the Microsoft Store is an abomination. But it's kind of like this switch on macOS. Look, allow apps downloaded from either App Store and identified developers. That means you have to have a certificate. Or App Store you can actually bypass this if you download something without a certificate, but you have to manually approve it. And if you put this in App Store and then locked it so that your kids or your users couldn't change it, that would be effectively S mode.

Steve: Yes.

Leo: That's Apple's version of S mode. And that's what I hope Microsoft does. I think that would be a good thing. I don't have a problem with that.

Steve: So his statement says, he claims, he says: "Since that time," that is, the time of the S on the Surface - maybe, oh, that's probably what "S" stood for, too, Surface and Store...

Leo: Right.

Steve: ...and Security and Stability and Safety and the American Way. Oh, no. Anyway, so since that time...

Leo: Microsoft always says it stands for nothing.

Steve: Okay. "Since that time," he says, "we've received great feedback from customers and partners on Windows 10 S. Customers love the security, the faster boot time..."

Leo: Oh, please.

Steve: It makes it faster, too.

Leo: No, it doesn't, though. It's Windows 10 Pro. It's the same.

Steve: Right. "Better battery life." Oh, it makes your battery life apparently longer.

Leo: Yes. See, if any of this were true, it might make it worthwhile. But as far as I can tell, it's not.

Steve: Right, right. And, he says, "...and consistent performance over time." What, because you can't load anything on it.

Leo: Because you can't put crap on it. Which is a real problem in the Windows world.

Steve: So he says: "Our partners have brought to market more than 20 devices with Windows 10 S enabled. We have also heard feedback that the naming was a bit confusing for our customers and partners." And of course we've done nothing to alleviate that confusion. And he says: "Based on that feedback we are simplifying" - oh, it stands for Simplifying - "the experience for our customers. Starting with the next update to Windows 10" - which I did a little bit of research. Apparently it's expected to be the spring, so maybe it stands for Spring, the Spring Creators Update, Windows 10 Spring Edition. Anyway, so with the Windows 10 Spring Creators Update, we believe. And he's saying: "Starting with the next update to Windows 10, coming soon, customers can choose to buy a new Windows 10 Home or Windows 10 Pro PC with S mode" - it's a mode - "enabled, and commercial customers will be able to deploy Windows 10 Enterprise with S mode enabled."

He said: "We expect the majority of customers to enjoy the many benefits" - okay, I said "many" - "the benefits of Windows 10 in S mode. If a customer does want to switch out of S mode, they will be able to do so," Leo, "at no charge" - you won't be charged for getting all the rest of Windows back - "regardless of edition. We expect to see new Windows 10 devices ship with S mode, available from our partners in the coming months, so check back here for updates."

Leo: I'm waiting for one, actually. It should come this week.

Steve: Uh-oh.

Leo: An HP.

Steve: Now, does that mean you can't turn it back on after you've been bad and turned it off?

Leo: You can, but it's not easy. You have to reinstall Windows 10 S mode.

Steve: Okay.

Leo: Because I did it on my Surface laptop. And then, by the way, I did it as an experiment because Paul wanted to know, so I bought the Surface laptop with 10 S on it, put up with 10 S for a while, but said, nah, I really want Chrome. And so I then flipped it off. It's very easy. It doesn't make many changes because frankly all it does is it says, go ahead, get stuff from the store, it's fine.

Steve: It loosens its grip.

Leo: It loosens its grip. It's very quick. And that was free. And then as an experiment we were curious, well, now that you've installed Chrome on it, if you go back to S mode, what happens to Chrome?

Steve: Ah, yes, I'm curious, too.

Leo: And that's why it takes so long. You have to download S mode, reinstall. And, yes, Chrome is gone.

Steve: Oh, it expunges your...

Leo: Expunges.

Steve: Okay.

Leo: And I think that's why it's more complicated to switch it off and on. I would love to see a switch - see, what would be great is switch it off, just like I can on the Mac, get Chrome, then switch it back on again.

Steve: Yup, and then leave me alone.

Leo: And then leave me alone; right? That would be ideal.

Steve: So protect me, but forgive my past sins.

Leo: Exactly.

Steve: Yes.

Leo: That would be nice. We don't know what they're going to do.

Steve: Oh, S for Sin mode. No.

Leo: Sin, right. Subsin. Sinless.

Steve: So for our listeners, something is coming with the Spring Creators Update that will sort of - it'll be interesting. He talks about being able to turn it off. And I guess there were some rumors that you'd have to pay an extra $49 to turn it off?

Leo: You used to have to.

Steve: Ah. You used to pay more?

Leo: Yeah, you've got to pay 50 bucks to go to Windows 10 Pro. So that was...

Steve: Oh, okay, Pro.

Leo: Yeah. Well, it is Pro, by the way. That may change.

Steve: Underneath the S there's Pro hiding.

Leo: It was always Pro, yeah. And the deal with the laptop, Surface laptop is you have till December, the end of the year, and you can get it for free. And after that it'll be 50 bucks. But I think they never really did that. I think they never really did that. It seemed like a good idea at the time.

Steve: Okay.

Leo: Anyway, I don't know. I've given that laptop to Megan. My guess is it's not still in S mode.

Steve: So many words start with S that it's just too much fun.

Leo: I know. But I am getting one of these new Windows on ARM machines because I'm very curious what they'll be like, running on a Qualcomm. Nothing wrong with Qualcomm chips, is there?

Steve: No, no, as far as we know. Although apparently that merger was blocked by our President.

Leo: Yeah, the President, which is really interesting.

Steve: Yeah, citing cybersecurity or national security.

Leo: Yeah, because Broadcom's in Singapore. Even though they were planning on repatriating to the U.S. We're in an old Broadcom facility, by the way. This studio was Broadcom.

Steve: Got great plumbing. Great plumbing.

Leo: It was so easy to get the servers in and everything, it was great. Hello, China. But, yeah, I thought that was very interesting. I don't know, he didn't give us any information about the justification except to say for security reasons. Singapore is a U.S. ally. Maybe he thinks Singapore is China. I don't know.

Steve: Okay, well, I'm just going to bite my tongue and move on.

Leo: Whatever. Hey, we're safer. We're all safer.

Steve: So we have an interesting collection of cryptocurrency / malicious mining / cryptojacking news. I got a kick from a note that I saw on the BleepingComputer site. Our friend Lawrence Abrams, who is the founder of BleepingComputer, which is an often mentioned on this podcast terrific site. They were early the main source of really good information when the very first cryptomalware, the file encryption stuff happened. I think they were the people who were on the top of that news. And then the so-called ransomware; right?

So I got a kick out of this because he said, I think it was late last week, he said: "It has been a pretty slow ransomware week, as most of the malware developers have started pushing crypto miners." So there you have it from the site that has been at the forefront of ransomware. And as we've been saying, as soon as it became possible to monetize attacks in some other way, like mining, that's what happened. So here's Larry saying it's been a pretty slow ransomware week as most of the malware developers have switched to crypto mining.

So in a different story there that sort of relates to that, they noted that Check Point Security did a study of the ranking of different malware. And in BleepingComputer's coverage they said: "Three in-browser cryptocurrency mining scripts ranked first, second, and fourth in Check Point's most active malware top 10." So across all malware, cryptocurrency mining is three of the top four and now outranks classic high-output malware distribution infrastructures such as spam botnets, malvertising, and exploit kit operations. That is, what it's hip to do now is get your mining address, your cryptocurrency address into as many other people's machines as possible, steal their computation cycles, and participate in a pool and turn it into cash flow.

So the number one, not surprisingly, the number one mining script is Coinhive. Number two is Crypto-Loot, that we have not yet encountered on this podcast. And number four is JSECoin, so no doubt JavaScript Electric coin. "These three are online services," they write, "that offer JavaScript libraries that website owners can embed on their sites and generate profit by using their visitors' CPU resources to mine Monero cryptocurrency."

And they note that, while all three are legitimate - and Coinhive, I've been often quoted actually from this podcast noting that I thought maybe they should really be held responsible because there's no way they don't know they are being abused. And we talked about it, in fact, the opt-in, the explicit opt-in face of Coinhive on a different domain and asking people, please don't block us. We're really going to make sure people opt in. We're taking responsibility for the opt-in ourselves so that nobody can use our script from this domain without us making sure the user knows that's happening.

That's really the only thing they could do at this point because they're now - they are a known domain that is being blocked by many of the adblockers. So in Check Point's case, the company said that its security products - and this is surprising, but I guess this is a function of Check Point's demographic, their customer demographic. So they said that their security products have detected cryptojacking across 42% of the organizations they protect, 42% of Check Point's customers that are contracted to get Check Point security protection. Forty-two percent have had cryptojacking mining detected, Coinhive, of course, in the lead with detections found on 20% of all of Check Point's customers, followed by Crypto-Loot in a close second place at 15% of all of Check Point's customers. So, I mean, it really is prevalent. It's amazing.

And how does the crypto mining deal with the adblockers? Well, not surprisingly because, again, the bad guys are as clever as the good guys, we have discussed the state-of-the-art means for botnets avoiding having their command-and-control server domains commandeered and taken over. And that's the so-called Domain Name Generation Algorithms, where even when the malware is reverse-engineered, what happens is that, based on time of day, which the networks are able to obtain from network time or the system they're on and verify it and so forth, I mean, it's easy to know when it is, they generate pseudorandom domain names, like a bunch of them, like all at once, like a hundred. And the idea is that, as we've discussed before, the white hats have to acquire all of those domains if they want to block them all because only one of them needs to respond to the botnet, which is happy to try them all.

So it's a very clever way of hugely upping the ante. And the fact is I think it was the Conficker worm that was one of the early botnets that did this. It is still in operation because it uses this technique, and it is prohibitively expensive for any group to constantly be having to preemptively register a large block of domains which are changing constantly. I mean, the bad guys, similar to like DDoS attacks, for which there's really no good defense, have similarly found a way of solving the, oh, this is the domain that this botnet uses, we'll just stomp on it and cut off communication.

Well, not surprisingly, this same technique has now come to cryptocurrency mining, where it is no longer the case that adblockers can simply block www.coinhive.com in order to block the scripts. Now the domains are being generated with algorithms based on time, and they are constantly changing with the mining script being sourced from domains which are unpredictable moving forward. So we're seeing a technique that was developed earlier for a different application, now being moved - and unfortunately was powerful and successful - now being used to thwart simple domain name-based adblocking. That doesn't work for keeping those sorts of scripts off of machines now.

The only thing I can see that is going to work is people starting to become much more aware than they are now about their CPU usage and utilization. I have for years, I'm looking at it right now, I have Task Man running in the tray minimized, showing the little square that's all green. But when my CPU starts getting used - there are four of them, it's a quad machine - it comes up so I can sort of see what's going on. And I would imagine at some point someone will create something that alerts someone if one particular process over some length of time is never letting go of the machine. Of course then the counterattack or the countermeasure for that would be for the miner to periodically stop mining and then go [humming], and then start up again in order to look more like the kind of process that you might expect to see. So it's a mess. And it's going to be with us for a long time because it makes money.

Speaking of which, two different groups, the Imperva crew and the SANS Security folks have been keeping an eye out for other mining ware. And the Imperva crew discovered new cryptojacking malware targeting Redis servers. Redis is a BSD-based caching server platform that offers services similar to memcache, but runs on a compatible piece of hardware and is unfortunately exploitable. And then there's also more cryptojacking going on targeting Windows-based servers. The SANS Security folks spotted a different campaign which targets the Apache Solr, which is a large, scalable, searchable, open source database system. And as usual, these are publicly accessible, in all cases, publicly accessible server platforms which are essentially - we can sort of think of them as, in several cases, they have memcached sort of services. But they're not being used for denial of service attacks, they're being used to mine cryptocurrency.

So again, the pressure that this creates to mint dollars or whatever denomination of currency is causing bad guys to come up with new ways of exploiting existing problems. And in every case they are exploiting previously patched, in some cases by a year or two, vulnerabilities which are well known, have been long fixed. But there are servers still operating that have not been updated; and the malware, which sort of ignored these services or servers until now, well, now there's a reason to get in there. So they're running their cryptocurrency mining on them.

Oh, and I had a couple stats also. In the case of the Redis servers, last month a different campaign than the one that was seen, but also against actually both Redis and OrientDB servers, netted nearly $1 million. So this is the other reason. As we know, the stronger the machine is, the more muscle its CPU has, the better the rate of cryptocurrency mining. But this demonstrates that whoever it was behind that campaign who took the time and trouble to find and infect those two classes of servers, the Redis and the OrientDB servers, during a month netted just shy of a million dollars.

So this is why we're seeing a lot of this happening and expect to in the future. It pays off. And although I don't have a dollar figure, the SANS researchers determined that around 1,777 infections of that Apache Solr had succeeded between the end of February and March 8, so February 28 to March 8, so essentially nine days, in order to find 1,777 Apache Solr instances and set up cryptocurrency mining there.

Leo: Apache Solr was the plugin used by Equifax that was unpatched that got them breached.

Steve: Ah ha ha, yup, makes sense. Cool. Also last week, exactly a week ago on March 6th, Microsoft detected a rapidly spreading cryptocurrency malware which successfully infected nearly half a million Windows machines within 12 hours. A week ago, Windows Defender suddenly detected more than 80,000 instances of a malware named "Dofoil," or also known as "Smoke Loader," which was dropping a cryptocurrency miner as its payload to mine "Electroneum" coins.

Leo: Sounds like a Marvel universe.

Steve: It does, doesn't it, yeah. Electroneum. That's a cool name, though. Go get me some Electroneum.

Leo: Of course, yeah, it's like Vibranium and Unobtainium.

Steve: Rub some Electroneum on that. That'll help it. Anyway, by the time Windows

Defender could respond, nearly 400,000 additional infections were in place, but they weren't local. They were rapidly spreading across Russia, Turkey, and Ukraine. The question which has so far gone unanswered by Microsoft is how such a large audience became infected in such a short period. We don't know that, but we do know that Defender shot off alarms, and then Microsoft was able to then update Defender and get it to help deal with that infection.

Okay. And one last cryptocurrency mining story, and then we'll take our final break. And that is, believe it or not, Leo, there's a currency miner in the Mac App Store, and Apple seems okay.

Leo: I saw this. Calendar 2, yeah.

Steve: Calendar 2, exactly.

Leo: Well, it's because it says - it's in the features. You can turn it on.

Steve: Correct, correct. And in fact the next page of the show notes I have a screenshot of Calendar 2's features. Now, it was a little controversial. Well, in fact the author has now said he's going to back out of that. But he thought, hey, I'm going to declare it upfront. I'm going to allow people to have all of the Calendar 2 features if they're willing, as he says, to unobtrusively generate cryptocurrency in the background.

Leo: That's actually a good deal. It's otherwise $18 to do that.

Steve: Yeah. So you could pay once for $18 to unlock all new and future calendar features, or you could go the pay-as-you-go route at $0.99 per month to subscribe to unlock all new and future calendar features, or just run with the totally free version, and you just get the basic stuff. So it was a little controversial because it was enabled, the free cryptocurrency mining version was what you got by default, without explicit notice.

So the critiques I saw suggested that, well, he really should have made it very clear, like asked users do you want to unlock all the features if we just do a little mining in the background? Also, it was supposed to be in the background. What the author now says is that there were problems with the library that they got from a third-party, whom he's blaming for these problems. It was only supposed to take 10 to 20% of the CPU. On the other hand, we know the more of CPU you take, the more you get. So for whatever reason, whether it was deliberate or not, it was using way more of the system, which of course brought the device, the Mac, apparently really slowed it down. And he was mining without permission.

And apparently, even if the user did not give permission, then due to some glitch in the library, we believe, it was causing a problem. I think that's all true because this ended up not going over well, and he's now said he's going to remove that. But again, as you said, Leo, if it were able to be unobtrusive, I guess maybe only mining while you're using the calendar, I'm assuming that's the case, it's not just running in the background all the time, and as an alternative to paying money, $18, maybe it makes sense.

Leo: Why not, yeah.

Steve: I think we're going to see more of this. To me, as I've said, as an alternative to web-based advertising, if it's efficient, that's the quest. And in fact, you know, I've got a really fun story. So, okay. Our last cryptocurrency story is just too funny. Well, fun. Get this. A French heater manufacturing company, and I don't know how to pronounce their name, Qarnot, Q-A-R-N-O-T?

Leo: Qarnot, I'd say Qarnot, yeah.

Steve: Qarnot is actually selling, or at least offering to sell, we don't really know how many people are doing this, a cryptocurrency mining-based space heater.

Leo: It's got an intuitive interface.

Steve: It does. It'll talk to your smartphone. It's got little LEDs.

Leo: It's the first crypto heater.

Steve: It's the QC-1, Crypto Heater QC-1.

Leo: It's noiseless.

Steve: Oh, yes, noiseless. No fans. No moving parts, so it uses...

Leo: Oh, a little pricey.

Steve: It's a little pricey, yes. It is 2,900 euros, which is at $3,600. It's able to hash at 60 megahashes per second.

Leo: Is that good?

Steve: Well, it makes Ethereum. And at current hash rates, difficulty, and Ethereum trading rate, that would be about $120 per month. So I don't know.

Leo: And they don't - electricity not included.

Steve: Right. And so it's not going to probably pay for itself. I mean, if you sit on it, your butt'll get warm. So on their website they say "Earn cryptocurrency while heating. The heat of your QC-1 is generated by two graphics cards" - they have a pair of Nitro Radeon RX 580s with 8GB of RAM, which as I said is able to solve Ethereum at 60 megahashes per second. They say "...two graphics cards embedded in the device and mining cryptocurrencies or blockchain transactions. While heating, you create money. You can watch in real-time how crypto markets are trending on your mobile app and on your QC-1 LEDs." It's got a little strip of LEDs.

And it says: "Noiseless high-end design. The QC-1 crypto heater" - it's hard to even say with a straight face. Yes, it's a crypto heater - "is the only one in the market [imagine that] to be perfectly noiseless." Maybe there's some noisy ones. "It doesn't embed any mobile part - no fans, no hard drives. This system developed by Qarnot is IP protected," meaning intellectual property, I'm sure. So they came up with the idea, and they're just going to sit on it. Well, or maybe not, depending upon how hot it gets.

"Intuitive interface, designed as a plug-and-mine device." You just plug it in and hook it to your WiFi and off it goes, maybe defraying the cost of heating your home a little bit. The problem is, again, certainly now where you're using a pair of strong GPUs, your efficiency makes sense. But again, $3,600 U.S. for a space heater? I don't know. Don't know if that quite cuts it. But give them an A for coming up with the idea.

Okay. I mentioned that tomorrow, March 14th, is the first day of the 2018 Pwn2Own conference up in Canada. This is its 12th annual competition. We've been covering them since 2007, so probably for as long as the podcast has been going since we're in our 12th year also. And we've always had a lot of fun with the Pwn2Own conferences, the results over the years. This year there's five categories: virtualization, web browsers, enterprise applications, servers, and Microsoft - that is one of the co-hosts along with VMware this year, a sponsor - the fifth category is a special Microsoft Windows Insider Preview Challenge category, where actually they figure out what the S for Windows 10 S is actually supposed to stand for. No, I'm kidding, I don't know.

The main backer, or the organizer, is Trend Micro's Zero Day Initiative, ZDI. And as I said, the partners for the event are Microsoft and VMware. As a consequence of them and other sponsors, up to $2 million in U.S. cash and prizes will be awarded.

Leo: Wow, that's huge.

Steve: It is huge. It's getting significant. Microsoft offers a Windows Insider Preview Challenge that tests their latest prerelease offerings combined with their configuration on their hardware, and the title "Master of Pwn" will be awarded to the team with the most points at the end of the contest. And so I thought there were a couple interesting things in the Zero Day Initiative announcement.

They said: "Since its inception in 2007, Pwn2Own has increased the challenge level at each new competition, and this year is no different," they wrote. "Web browsers return as a target" - yes, it's a fan favorite, folks - "as do virtual machine guest-to-host escapes. Enterprise applications remain as targets for this year, and for 2018 Outlook [oh, boy] makes its Pwn2Own debut." In years past, Outlook would have just been shredded. It'll be interesting to see how it fares.

They said: "Our virtualization category grows by two as Oracle becomes a target, and the Windows Insider Preview Challenge includes brand new targets for their virtualization-based security stack. Server targets expand this year, as well. Apache was included in last year's event and is joined this year by NGINX, OpenSSL" - oh, this is going to be a really interesting competition - "and Windows SMB server [oops]. Over the years we've seen some groundbreaking research demonstrated, so we can't wait to see," they write, "what contestants bring this year." However, this year there will be no hacking entrants from the previous year's winning country.

Leo: Oh.

Steve: China's government has decided to keep its security researchers home. And as a new policy, Chinese researchers will not be attending remote events. Now, my first thought was that this might be to prevent them from being detained as we know Marcus Hutchins was after last summer's Black Hat conference in Las Vegas, where he was nabbed at the McLaren Airport as he was attempting to return to London, I think it was. So Brian Gorenc, who's the director of Trend Micro's Zero Day Initiative, said: "There have been regulatory changes in some countries that no longer allow participation in global exploit contests such as Pwn2Own and Capture the Flag competitions." And although he didn't explicitly in this quote mention it, he was referring specifically to China. So there will be no Chinese research teams at Pwn2Own this year, which will likely be felt since for the last several years the Chinese teams dominated the competition.

So I think it's sort of sad because they've got excellent hacking skills. Although I suppose for the people they were competing against, having China held back sort of shifts the balance. So it'll be interesting to see who the winners are when China is not present. Prior year Chinese winners were contacted and asked for comment, but none would remark other than to indicate that they would not be attending the competitions.

And in another little bit of Chinese-related news, a U.S. threat intelligence firm known as Recorded Future has spotted that the Chinese equivalent of our CVE database - we often talk about CVE, that's the Common Vulnerabilities and Exposures database which is located at cve.mitre, M-I-T-R-E, dot org. China has one known as the CNNVD, the Chinese National Vulnerabilities Database, which is similarly meant to be open kimono and available and accessible. Anyway, this firm has discovered that China has been retroactively manipulating the data in their own Chinese national vulnerabilities database, basically playing fast and loose with the facts and altering the database to conceal the influence of the Chinese Ministry of State Security.

They produced a report with four key judgments. They said the CNNVD altered the original publication dates in its public database for at least 267 vulnerabilities that they identified as statistical outliers in their research which was published in November of 2017. They wrote: "We assessed in November that CNNVD had a formal vulnerability evaluation process in which high-threat CVEs were evaluated for their operational utility by the MSS [Ministry of State Security] before publication, and that the publication lag was one way to identify vulnerabilities that the MSS was likely considering for use in offensive cyber operations."

Now, in all fairness, it's hard to imagine that our own intelligence services in the U.S. aren't - who knows what manipulation is going on? There has been, we have had some reporting of this notion that some of these vulnerabilities are being disclosed to the U.S. government before being made public. So basically a similar operation.

Leo: And vice versa, Intel told Lenovo before they told the U.S. government.

Steve: Exactly. That was quite controversial, yes. And although they say here that CNNVD's outright manipulation of these dates implicitly confirmed this Recorded Future's assessment. They said: "By retroactively changing the original publication dates on these statistical outliers, CNNVD attempted to hide the evidence of this evaluation process, obfuscate which vulnerabilities the MSS may be utilizing, and limit the methods researchers can use to anticipate Chinese APT [Advanced Persistent Threat] behavior." And then they conclude, saying: "This large-scale manipulation of vulnerability data undermines trust in the CNNVD process and could compromise security operations relying solely on the CNNVD for that information." Although, again, our own CVE database is similarly public and probably contains virtually all of the same information.

I did want to note, just in passing, that last Friday U.S. District Judge Lucy Koh denied in part a motion by Verizon, which is the owner of Yahoo, to dismiss a case brought by plaintiffs who are suing Yahoo for failing to adequately protect their users' security and neglecting to respond to known and dangerous vulnerabilities in a timely manner. We know from our previous reporting, and widely covered in the industry, that December before last, December 2016, Yahoo first disclosed the truth of the 2013 breach which they said compromised the credentials of a billion of their users. And then just last October that assessment was revised to three billion impacted accounts.

Judge Koh wrote in her ruling on Friday: "Plaintiffs explain that, had they known about the inadequacy of these security measures, they would have taken measures to protect themselves. Plaintiffs' allegations are sufficient to show that they would have behaved differently had Defendants disclosed the security weaknesses of the Yahoo Mail system." Which it now appears was known to Yahoo. And at the time we talked about how Yahoo! just decided not to follow the urgent requests of their own security staff to fix problems, but rather spend their resources elsewhere. So it'll be interesting to see how this settles out. It is significant, I think, that Yahoo! is being held to account because other companies are certainly keeping an eye on this.

And finally, this stunned me, Leo. I was very surprised by this. I have a picture of, I guess it's a doughnut chart. We used to just have pie charts. Now we have doughnut charts, for some reason, I don't know why. But the news is, according to a report that McAfee just released yesterday, McAfee analyzed the source of email spam Internet-wide. And I was amazed by the distribution. Two botnets, one called Necurs, N-E-C-U-R-S, and Gamut, G-A-M-U-T. Those two botnets alone account for 97% of the Internet's spam, email spam. I'm amazed. And so this doughnut chart, this Necurs is by far the biggest. It's generating 60% by volume in the fourth quarter of last year of all of the spam email. And Gamut is at 37%. So together 97%. And then the other three are Lethic, Darkmailer, and then random Others, each with only 1% respectively.

So in McAfee's report they said: "For most of these months, Necurs has spent its time churning out [what they called] 'lonely girl' spam lures for adult websites, pump-and-dump schemes, and delivering ransomware payloads. Overall, nearly two out of three spam emails sent in the last quarter of 2017 were sent from the infrastructure of this mammoth," writes McAfee, "botnet. Second on the list was the Gamut botnet, also built on Windows machines infected with malware that hijacks Windows to send out spam. Gamut, while smaller in size when compared to Necurs," they write, "had previously been more active in the third quarter," sending more spam than Necurs. However, they reversed positions. In Q4, Gamut activity went down, but the botnet still accounted for 37% of all email spam, compared to Necurs' 60%.

They wrote: "Most of Gamut's email subjects were related to job offer-themed phishing and money mule recruitment, [which is] tricking people to buy products with stolen money and sending the products to crooks; relaying money from hijacked bank accounts" to malicious accounts and so forth. But I'm just amazed. Two botnets, two massive Windows botnets are generating 97% of the spam on the Internet. Wow.

Leo: That is so amazing.

Steve: Huh? Go ahead.

Leo: It's amazing, yeah.

Steve: Yes, isn't that a surprising distribution? I'm just stunned. It's like they put everybody else out of business in the spam business.

Leo: Yeah.

Steve: I got a nice note at the end of February, on the 27th, from Eric in Wisconsin. The subject was "Yet Another SRSTD Story." And then he says, his first line in the email is "SRSTD = SpinRiteSavesTheDay."

He said: "Steve, first of all, THANK YOU [all caps] for your great SpinRite product." He says: "I've been meaning to pass along a story of yet another of SpinRite's successes for me for quite some time. I have often thought this may be just another run-of-the-mill testimonial of how SpinRite has saved the day, but I thought I would share it anyway and let you be the judge." And of course I'm always happy to have those and to be able to share them with our listeners.

He says: "Besides, I had to say thanks. A dentist's office was referred to me for support when their <insert well-known brand here> NAS device" - probably a Synology, just based on what he says later - "an NAS device was indicating multiple drive failures" - sounds like they'd neglected it for quite a while - "and they could no longer access any of their data for the office because the NAS would no longer boot." He says: "Of course I asked about their current backup process and was shocked [he says in quotes, parens] ('not really') to hear they had been meaning to address that hole in their process for some time." So no backups.

"This particular NAS was bootable from its USB port," he said, "so I booted it into SpinRite and let SpinRite go to work. By now, everyone knows that a happy ending usually ensues, and that is an understatement in this case. SpinRite plowed through the entire array and upon reboot was happily serving data again. For safekeeping, I was able to get a copy of their never-backed-up data off the NAS. Needless to say, they were ecstatic about the recovery and also now have a good local and offsite backup process in place." Well, clearly Eric has been listening to the podcast and knows the importance of doing that and obviously is up to speed on how to recover data and then keep it safe.

He says: "I did recommend to them to go purchase their own copy of SpinRite to express their satisfaction, and I believe they have done so. Thank you again for ALL [in all caps] that you do. Look forward to listening to about 350 more SN episodes." Yes, up until we finish with three digits. He says: "Keep up the good work. Eric in Wisconsin." So Eric, thank you for sharing that with me and allowing me to do so with our listeners.

And so finally, what was going to be the topic of today's podcast was to follow up on last week's podcast, where we introduced our listeners to what had just happened in the world, which was the discovery that there were tens of thousands of open, exposed to the public Internet, servers running memcached, a daemon, a server process, which was listening on port 11211 and would both receive data over that port for local caching and then, based on a tag that that data was tagged to, when later asked for it would send it back out. The problem is since it was bound by default to both UDP and TCP - which, by the way has been changed in the updated version. Now it's only bound to TCP, which is a minimal change. But on the other hand, we know how few of these are probably going to be updated over time. So these big available caches are going to be sitting on the public Internet for quite a while.

Now, for those who do update, then it will no longer be publicly exposing a UDP port. That's important because it is trivial to change the source IP of a UDP packet in order to cause the memcache server to believe somebody else asked for this big blob of data. As I mentioned at the top of the show, it is possible to create a bandwidth amplification of 51,000x. So just crazy.

And in fact I didn't mention last week that there have been ransom demands being made of the sites that are being DDoSed, and the ransom demand is in the data that the site is being DDoSed with. If you're being flooded with that much traffic, it's unlikely that many of those demands actually make it through your pipe, like to you, because it's going to crash upstream routers and just wreak havoc on your system because these attacks are so massive. But I think it sounds like the attackers had a little bit of a sense of humor.

What we have seen, as the Picture of the Week at the top of the show notes showed, was predictably a massive increase in these attacks. In just the last 10 days - and this was a report that ended several days ago. So essentially in the week that followed the emergence of this, over 15,000 DDoS attacks based on this memcached service were used to attack 7,100 sites over that period of time. Also since then the concept of course went wild. There is now publicly available, freely available, downloadable attack code.

So although this was never a difficult attack to launch, and I'm sure a lot of coders who are capable of coding their own attack tools did so immediately, now even that's not necessary. There is downloadable turnkey DrDOS - Distributed Reflection Denial Of Service - attack code available where you simply run the tool. You put in the target, the IP address of the target you wish to attack. That tool goes out to Shodan and looks up, gets the current list of known to Shodan memcached service IPs and begins filling the cache and bouncing UDP packets requesting that that cache be dumped on whatever victim was targeted. So that no longer requires any coding. That stuff is freely available on the Internet.

And the last little bit of news is that there has been some attempt at considering mitigation. There are two other commands, other than "here's a tag, send me back the blob." There is an "empty your cache" command, and there is a "shutdown" command. I haven't seen anybody suggesting the use of the shutdown command. And of course it would bring, if the memcached service was in active use by the site, having it shut down would bring it to the attention of the site owners pretty quickly.

But there is a "flush the cache." And I've seen some suggestions that, if a site was under attack, that the "flush the cache" command could be sent proactively back to their array of sites that are attacking it. Unfortunately, that's like 17,000 different services, at least. And the moment you start flushing, then the other sites are still attacking you, and the flushed cache can begin to be refilled. So the problem is - not that there's any practical solution to these attacks. It's that we have publicly exposed caches now available on the

Internet that are by default bound to both UDP and TCP, UDP trivial to spoof. And so as would have been predicted, massive attacks are now underway leveraging this.

The only thing that I've seen is that some providers, some big providers like Cloudflare are throttling all traffic across their border across port 11211. That's the well-known port, as I mentioned, where this service is running. So by throttling that traffic, at least that prevents both outgoing requests for cached data to remote public servers and prevents the data from getting through their border. On the other hand, the bandwidth flood will still hit their border and cause a massive attack.

So at this point all we can do is keep an eye on this and see what the trend is moving forward. This isn't crypto mining. It's not making anyone any money, unless ransoms are being paid. It does make DDoS, very, very potent DDoS attacks, far more powerful and easy to create than they were before. So maybe it shifts the balance. Or maybe this'll die down after a while, and people will go back to trying to install cryptocurrency mining, which is the new trend on the Internet, making some money that you can exchange for real dollars. We'll see.

But the big story, of course, is AMD's very, very short notice of some serious problems. I'm sure we will have much more to say about that on the podcast next week.

Leo: You bet. Or, as some speculate, it's just some company trying to dump AMD stock since they didn't publish exploits; right? I mean, is that conceivable? Awful lot of work to put in, a lot of technical detail.

Steve: Yeah. And as I said, I was tempted to push this off to next week until I read through their research, which is very compelling.

Leo: Yeah, okay.

Steve: So, yeah, I think that's not the case. And does anyone know what has happened to AMD stock in the meantime?

Leo: You know, it's funny, not much.

Steve: I'm not an owner of any tech stock.

Leo: No, we don't have tech stocks.

Steve: Of any stock at all.

Leo: You don't have stock at all.

Steve: No.

Leo: I mean, you must have mutual funds or something. You have investments.

Steve: I'm a municipal bond owner.

Leo: Oh, there you go. That's a good [crosstalk].

Steve: That's my solution. Long time ago, back when the interest rates were much higher.

Leo: Let me just check the AMD stock price. Yeah, it's up 1%.

Steve: Yeah, well, it'll be interesting. Again, I don't know if the news has hit mainstream. Maybe there's already some vulnerability fatigue from Spectre and Meltdown, where it's like, okay, well, another one of these big problems. We dealt with the ones at the beginning of the year, we'll deal with this one. Who knows. Anyway, it looked absolutely credible to me. I'm sure I'll have much more mature coverage, because much more will be known, by a week from now.

Leo: There was a bit of a selloff when the show began, but it's flattened out now in the afterhours trading. Yeah, I don't - it's hard to say that the stock has been hurt much.

Steve: Yeah.

Leo: It was going down for a while. All right, Steve. Somebody said Steve owns all of Fresno. He owns all their convertible debt, yes. Thank you, Steve. Steve Gibson is at GRC.com. That's his website, and he tweets @SGgrc. You'll find him on the Twitter there. You can leave long DMs for him if you want, if you have comments. I left you a tweet, a funny article I read that might make a good Photo of the Week next week.

Steve: Oh, cool.

Leo: But Steve always likes more input like that. You go to GRC.com you'll find SpinRite, that fine tool for dentists and hard drive users everywhere. All you need to do is go to GRC.com, pick up a copy, and you'll be ready to test and recover hard drives to your little heart's content. There's plenty of other stuff there, too, lots of freebies, including of course this show, 64Kb audio and fine handwritten transcripts from Elaine Farris, all there. Searchable, too, all 654 episodes.

You can also find audio and video at our site, TWiT.tv/sn, or subscribe. You'll get your show as soon as it's ready on your favorite podcast platform. Security Now! has been around for a little while, so we're on most, I think almost every program has us. You can even ask your voice-activated device to listen to Security Now!, and most of the time that works pretty well. I think that's it for the day and for the week. But we will see you, we will convene against next Tuesday, 1:30 Pacific, 4:30 Eastern, 20:30 UTC.

Steve: Oh, I'm sure we'll have more news, my friend.

Leo: Thank you, Steve. See you then.

Steve: Thanks, Leo.