Security Now! #655 - 03-20-18

Pwn2Own 2018

This week on Security Now!

This week we discuss the aftermath of CTS Labs' abrupt disclosure of flaws in AMD's outsourced chipsets, Intel's plans for the future and their recent microcode update news, several of Microsoft's recent announcements and actions, the importance of testing... in this case VPNs; the first self-driving automobile pedestrian death, a SQRL update, a bit of closing the loop feedback with our listeners, and a look a the outcome of last week's annual Pwn2Own hacking competition.

Our Picture of the Week

Richard Zhu (@fluorescence) won the contest by gaining 12 points for hacking Edge and Firefox. Zhu took home $120,000 of the $267,000 total prize money awarded at the event.

Security News

RyzenFall Update
Controversy surrounding CTS Labs (lack of) disclosure courtesy delay.
Hyperbole Swirls Around AMD Processor Security Threat
AMD Flaws Pose No Immediate Risk of Exploitation, Says Independent Reviewer

Dan Guido, security researcher and CEO of Trail of Bits, independently verified the reports of the vulnerabilities and exploit code on Twitter: "Regardless of the hype around the release, the bugs are real, accurately described in their technical report (which is not public afaik), and their exploit code works." — Dan Guido (@dguido) March 13, 2018

On the AMD Flaws website: A Letter from our CTO
https://amdflaws.com/
https://safefirmware.com/CTO+Letter.pdf

Disclaimer

I am a technical guy, I love building things, and researching things. I do not like the whole world of PR, it is too messy for me. I have written this letter in my own language, without PR proofing, please forgive me if there are any grammatical errors, or not written according to correct writing standards.

History of the publication

We have started researching ASMedia chips about a year ago. After researching for some time, we have found manufacturer backdoors inside the chip which give you full control over the chips (ASM1042, ASM1142, ASM1143). We wanted to go public with the findings, but then saw that AMD have outsourced their chipset to ASMedia. So we decided to check the state of AMD, we bought a Ryzen computer, and whimsically ran our exploit PoC, and it just worked out of the box. Full Read/Write/Execute on the AMD Chipset, as is – no modifications. To be honest, we were a bit shocked by it, how they have not removed the backdoors when integrating ASMedia IP into their chipset is beyond me. So then we said, ok – what on earth is going on in AMD, and started researching AMD.

It took time to set-up the working environment to start communication with the AMD Secure processor, but after reaching a full working setup and understanding of the architecture – we started finding vulnerabilities. One, and another and another. And not complex, crazy logical bugs, but basic mistakes – like screwing up the digital signatures mechanism. At that point, about once a week we found a new vulnerability, not in one specific section, but across different sections and regions of the chips. It's just filled with so many vulnerabilities that you just have to point, research, and you'll find something (obviously a personal opinion).

After that we decided we have to go public with this. I honestly think it's hard to believe we're the only group in the world who has these vulnerabilities, considering who are the actors in the world today, and us being a small group of 6 researchers.

Responsible Disclosure

I know this is an extremely heated topic for debate, where everyone has a strong opinion. Unfortunately, I also have a strong opinion on this topic.

I think that the current structure of "Responsible Disclosure" has a very serious problem. If a researcher finds a vulnerability, this model suggests that the researcher and the vendor work together to build mitigations, with some time limit (30/45/90 days), at the end of which the researcher will go out with the vulnerabilities. The time limit is meant to hasten the vendor to fix the issues.

The main problem in my eyes with this model is that during these 30/45/90 days, it's up to the vendor if it wants to alert the customers that there is a problem. And as far as I've seen, it is extremely rare that the vendor will come out ahead of time notifying the customers – "We have problems that put you at risk, we're working on it". Almost always it's post-factum – "We had problems, here's the patch – no need to worry".

The second problem is - if the vendor doesn't fix it in time – what then? The researcher goes public? With the technical details and exploits? Putting customers at risk? How we have accepted this mode of operation is beyond me, that researchers advertise at the end of the time limit the technical details of the vulnerabilities "because" the vendor didn't respond. Why should the customers pay for the vendor's lack of actions. I understand – this is the model today and people follow suit, but I think we can do better.

I think that a better way, would be to notify the public on day 0 that there are vulnerabilities and what is the impact. To notify the public and the vendor together. And not to disclose the actual technical details ever unless it's already fixed. To put the full public pressure on the vendor from the get go, but to never put customers at risk. This model has a huge problem; how can you convince the public you are telling the truth without the technical details. And we have been paying that price of disbelief in the past 24h. The solution we came up with is a third party validation, like the one we did with Dan from trailofbits. In retrospect, we would have done this with 5 third party validators to remove any doubts. A lesson for next time.

I know there are many questions, and a whole lot of confusion. We are trying our best to answer reporters, update our site with Q&A, and clarify what's going on. So far the media focus was on CTS, and I think I understand this, but very soon we will have to deal with the fact that a huge company with products spread throughout millions of computers in the world, is riddled with so many problems that it's unclear how to even address this.

If you have any technical questions, please contact me at ilialuk@cts-labs.com, I'll try to answer as many questions as I can.

Meanwhile: Intel Details CPU 'Virtual Fences' Fix As Safeguard Against Spectre, Meltdown Flaws
https://newsroom.intel.com/editorials/advancing-security-silicon-level/

An Editorial by Intel's CEO Brian Krzanich posted last week on March 15th was titled: "Advancing Security at the Silicon Level - Hardware-based Protection Coming to Data Center and PC Products Later this Year"

However, it was very scant on details.

Speaking to Intel's so-called "Security First" pledge, Brian wrote:

Today, I want to provide several updates that show continued progress to fulfill that pledge. First, we have now released microcode updates for 100 percent of Intel products launched in the past five years that require protection against the side-channel method vulnerabilities discovered by Google. As part of this, I want to recognize and express my appreciation to all of the industry partners who worked closely with us to develop and test these updates, and make sure they were ready for production.

With these updates now available, I encourage everyone to make sure they are always keeping their systems up-to-date. It's one of the easiest ways to stay protected. I also want to take the opportunity to share more details of what we are doing at the hardware level to protect against these vulnerabilities in the future. This was something I committed to during our most recent earnings call.

While Variant 1 will continue to be addressed via software mitigations, we are making changes to our hardware design to further address the other two. We have redesigned parts of the processor to introduce new levels of protection through partitioning that will protect against both Variants 2 and 3. Think of this partitioning as additional "protective walls" between applications and user privilege levels to create an obstacle for bad actors.

These changes will begin with our next-generation Intel® Xeon® Scalable processors (code-named Cascade Lake) as well as 8th Generation Intel® Core™ processors expected to ship in the second half of 2018. As we bring these new products to market, ensuring that they deliver the performance improvements people expect from us is critical. Our goal is to offer not only the best performance, but also the best secure performance.

KB4090007: Intel microcode updates
https://support.microsoft.com/en-us/help/4090007/intel-microcode-updates
Eight CPUIDs spanning Skylake, Kaby Lake and Coffee Lake.

Last Patch Tuesday Spectre & Meltdown updates:
We didn't discuss the nice Windows improvements which were part of this month's Patch Tuesday.

In March Microsoft released updates for its x86 version of Windows 10 and microcode updates for devices running the Windows 10 Fall Creators Updates and Intel's sixth-gen Skylake processors.

Most recently, the company offered new updates against Meltdown and Spectre with new releases on this month's Patch Tuesday for PCs running x86 versions of Windows 7 and 8.1 as well as Server 2008 and 2012.

So... GRC's InSpectre might be showing you some better news. :)

"Speculative Execution Bounty Launch"
https://blogs.technet.microsoft.com/msrc/2018/03/14/speculative-execution-bounty-launch/

Remember that Intel beefed up their own bug bounty program, valid through December 31st, which will payout up to a quarter of a million dollars for reports of new similarly-serious side-channel bugs. Bugs rated with severities of 9.0 and 10.0 on the (CVSS) Common Vulnerability Scoring System scale will pay out up to $250,000. Vulnerabilities rated between 7.0 and 8.9 will carry a bounty of as much as $100,000. Below the 7.0 threshold, awards max out at $20,000.

... and now Microsoft has duplicated this program:

In other words: If anyone else is able to find anything wrong like Spectre or Meltdown we're going to pay for that information... and IF it's something =NEW= it'll be worth a quarter of a million dollars to you!

"Today, Microsoft is announcing the launch of a limited-time bounty program for speculative execution side channel vulnerabilities. This new class of vulnerabilities was disclosed in January 2018 and represented a major advancement in the research in this field. In recognition of that threat environment change, we are launching a bounty program to encourage research into the new class of vulnerability and the mitigations Microsoft has put in place to help mitigate this class of issues."

Quick Facts:
- o Bounty Duration: Open until December 31, 2018
- o Full Details: Speculative Execution Bounty Program
- o Bounty Terms: Standard terms and conditions apply

Tier 1: New categories of speculative execution attacks: Up to $250,000
Tier 2: Azure speculative execution mitigation bypass: Up to $200,000
Tier 3: Windows speculative execution mitigation bypass: Up to $200,000
Tier 4: Instance of a known speculative execution vulnerability (such as CVE-2017-5753) in Windows 10 or Microsoft Edge. This vulnerability must enable the disclosure of sensitive information across a trust boundary Up to $25,000

Future major Windows updates will take longer to install, but will spend less time "spinning."
It should now take about half an hour during the reboot phase.
https://arstechnica.com/gadgets/2018/03/microsoft-promises-less-downtime-for-installing-major-windows-updates/

Microsoft estimates that the Creators Update, released almost a year ago, would take about 82 minutes on average during the "roller coaster dots" offline phase.

Improvements made in this past Fall Creators Update cut that down to about 51 minutes.

And the next update (which we've been provisionally calling the Spring Creator's Update) hopes to reduce this "system unusable" time further still... to just half an hour.

It's not that the upgrade process itself has been getting faster, it's the Microsoft has been working to migrate as much work as possible into an "online phase" during which the system can still be used while Windows is working in the background to get things as ready as possible for the still-necessary online phase.

Microsoft has indicated that this decreased interruption =does= come at some cost since the online portion will now take much longer... that that phase is performed at a low priority, so it should not, in general, hinder the use of the computer.

The importance of testing: 3 Major VPNs services were tested.
All three of those tested were leaking their client's IP address.
https://www.vpnmentor.com/blog/vpn-leaks-found-3-major-vpns-3-tested/

Hotspot Shield, PureVPN, and Zenmate VPN all suffer from readily exploitable IP leaks.

Hotspot Shield responded immediately and responsibly and fixed their problems.

PureVPN and Zenmate VPN... did not. They are currently vulnerable so details are not being disclosed.

VpnMentor hired a team of three external ethical hackers to find vulnerabilities in three random popular VPNs.

- One hacker wants to keep his identity completely private, but the other two are known as "File Descriptor" and Paulos Yibelo.

- "File Descriptor" is a reputable, ethical hacker working for Cure53, the company hired by TunnelBear to identify and fix issues with their VPN applications, and one of the leading companies in security research.

- Paulos Yibelo, managed the VPN testing team and is a reputable application security researcher. He has found vulnerabilities in popular VPNs and published them in the past. His work was mentioned in ZDNet, SlashDot, and other media sources.

Leaks Summaries:

The leaks would allow governments, hostile organizations, or individuals to identify the actual IP address of a user, even with the use of the VPNs.

Zenmate's leak was somewhat minor compared to the two other VPNs.

We believe that most other VPNs suffer from similar issues, so the fast response of Hotspot Shield is something we think is worth commending. We felt that they worked with our research team in a fast and serious manner and that they care for their users. They took our research as help for improvement rather than criticism.

Since it took longer for ZenMate and PureVPN to respond to us, we are only sharing information about the vulnerabilities that were found and patched in HotSpot Shield. We advise users of PureVPN and Zenmate to be wary of the leaks they may face and check with their VPN providers for an immediate fix.

Hotspot Shield details:

CVE-2018-7879: Hijack all traffic
The Proxy Auto-Config (PAC) script used in Hotspot Shield's Chrome extension allows for trivial traffic redirection: It looks for the query parameter act=afProxyServerPing, and if found, it routes all traffic to the proxy hostname provided by the server parameter.

Any malicious adversary could simply cause a victim to visit a link with those parameters, and all traffic will then be rerouted through an attacker's proxy server. (Whoops!)

CVE-2018-7878: DNS leak
DNS queries were not being routed through the VPN tunnel, they were being made by the system's local DNS... so any website is able to determine someone's source IP.

https://www.dnsleaktest.com/ reveals your DNS servers... and those servers know the querying IP.

CVE-2018-7880 IP leak

We observed the following PAC script:

let whiteList = /localhost|accounts\.google|google\-analytics\.com|chrome\-signin|freegeoip\.net|event\.shelljacket|chrome\.google|box\.anchorfree|googleapis|127\.0\.0\.1|hsselite|firebaseio|amazonaws\.com|shelljacket\.us|coloredsand\.us|ratehike\.us|pixel\.quantserve\.com|googleusercontent\.com|easylist\-downloads\.adblockplus\.org|hotspotshield|get\.betternet\.co|betternet\.co|support\.hotspotshield\.com|geo\.mydati\.com|control\.kochava\.com/;if(isPlainHostName(host) || shExpMatch(host, '*.local') || isInNet(ip, '10.0.0.0', '255.0.0.0') || isInNet(ip, '172.16.0.0', '255.240.0.0') || isInNet(ip, '192.168.0.0', '255.255.0.0') || isInNet(ip, '173.37.0.0', '255.255.0.0') || isInNet(ip, '127.0.0.0', '255.255.255.0') || !url.match(/^https?/) || whiteList.test(host) || url.indexOf('type=a1fproxyspeedtest') != -1) return 'DIRECT';

What we found is that the whitelist for DIRECT connection is just too loose.

Here are two examples we found:

> Any domain with localhost will bypass the proxy, e.g. localhost.foo.bar.com
> Any URL with type=a1fproxyspeedtest will bypass the proxy

Security Now takeaway: is that just using a VPN, especially with browser integration east-of-use features, is not automatically actually secure just because SOME of the user's traffic is encrypted.

We've had our first self-driving car involved pedestrian death.
Initial reports appear to suggest that the car's software was likely not at fault.

An Uber R&D automobile which was in fully autonomous driving mode -- with a backup driver -- hit and killed (she died not long after in a Tempe, AZ hospital) a 49 year old pedestrian who was walking her bicycle across the road at around 10pm yesterday night. This is the first pedestrian death.

Uber has suspended its driverless car testing program nationwide.

In their coverage of this, ArsTechnica reports: If the death was caused by an error by Uber's driverless car software, it would be the first such death in the United States.

However, Tempe police Sgt. Ronald Elcock said the Uber vehicle was traveling around 40 mph in a 35 mph zone when it hit Helzberg as she stepped suddenly, mid-block, out onto the street. Neither she nor the backup driver showed signs of impairment. Elcock said: "As soon as she walked into the lane of traffic, she was struck by the vehicle."

Tempe's Police Chief Sylvia Moir told the San Francisco Chronicle that the Uber had a forward-facing video recorder, which showed the woman was walking a bike at about 10 p.m. and moved into traffic from a dark center median. She said: "It's very clear it would have been difficult to avoid this collision in any kind of mode."

Moir added that "The backup driver said it was like a flash, the person walked out in front of them. His first alert to the collision was the sound of the collision."

An error in Tesla's Autopilot software contributed to the death of Tesla driver Joshua Brown in 2016, but that crash happened only after the driver ignored repeated warnings to put his hands back on the steering wheel. And Tesla has emphasized that Autopilot is a driver assistance feature, not a replacement for a human driver.

Uber, in contrast, is trying to build fully autonomous vehicles. There are hundreds of such vehicles being tested around the United States by Waymo, Cruise, Uber, and other companies. But until now, there have been few serious accidents caused by this testing and no fatalities.

Miscellany

SQRL Update:

- The Sensitive Data Monitor
- Screen darkening / Sandboxie / Keyboard Focus
- SQRL's own secure update system.

SpinRite

Anonymous Sender <anon@grc.com>
Location: Somewhere
Subject: Spinrite story in 652 could have been social engineering!
Date: 07 Mar 2018 21:01:24
:
Hi Steve. I just caught SN 652 listened to the SpinRite story.

Customer?: "I lost my copy of SpinRite. I am incapable of record keeping. I have no backups. I'm an unorganized person. BTW, the email I bought SpinRite with no longer works (wink!). Here's my new email address. Can I have a new copy?"

GRC Support: "No problem. We don't need to verify anything. We trust you. Here's a link to the download. Feel free to change your(?) contact information. We'll just email you the secret pin in plaintext."

Customer?: "Really? Yoink! If anyone asks, I was never here, and no, I am not wanted in 5 states for fraud."

How was this different from someone calling your cell carrier or your bank and claiming to be you?

Bank: "You cannot login using your email that you registered with? You don't remember your secret pin? No problem. We'll just send you the secret pin to your new email."

This could have been social engineering! It is often said the customer is always right, but what if they aren't a customer?

Looking forward to the end of SQRL and the beginning of SpinRite 6.1.

Closing The Loop

Kari Mattsson @KariMattsson
@SGgrc Any comments on this?
EV certificates under DigiCert root CA w/o human intervention.
https://trustcubes.com/ev-ssl/
Shouldn't EV certs be more scrutinized?

Itinerant Engineer @CallMeImperiale
@SGgrc @SwiftOnSecurity
"Perfect Forward Secrecy" ...is the... "I guess ya just had t'be there..." of encryption.

Pwn2Own 2018

China was sorely missed.

China's hackers routinely win, sweeping the board, notably, the Tencent and Keen teams who we've covered in past years' competitions. In last year's competition the top five winners were from China, with three of them from Tencent.

But this year, as expected, Pwn2Own was without a single Chinese entrant... and it showed.

There are potentially many bugs -- and cash prizes -- to be had. Last year, for the event's ten-year anniversary, the Zero Day Initiative awarded a total of $833,000 to white hat hackers, exposing 51 different zero-day bugs. And most of those were found by Chinese researchers.

But this year, despite pulling together a purse of $2 million, this years hackers took home only $267,000 of that total $2 million prize pool despite the target-rich environment which included Oracle's VirtualBox, VMware Workstation, Microsoft Hyper-V Client, Chrome, Safari, Edge, Firefox, Adobe Reader, Microsoft Office 365 ProPlus, Microsoft Outlook, NGINX, Apache Web Server, Microsoft Windows SMB, and OpenSSL. And where the largest prize was $250,000 for a sandbox escape to hypervisor or kernel for Microsoft's Hyper-V virtual machine client.

To top it off, several non-Chinese competitors dropped out at the last minute after the immediately previous patch Tuesday closed the holes they had been planning to exploit, demonstrate... and cash-in on.

The Pwn2Own winner was Richard Zhu who won the contest with 12 points for hacking Edge and Firefox. Zhu took home $120,000 of the $267,000 total prize money awarded at the event. Each researcher also got to keep the laptop they tried their exploits on.

Leo… please read in the voice of a fight announcer or gymnastics competition judge...

> Day One:
>
> The first day of Pwn2Own 2018 has come to a close, and so far, we've awarded $162,000 USD and 16 points towards Master of Pwn. Today saw 2 successful attempts, 1 partial success, and 1 failure. In total, we purchased 3 Apple bugs, 2 Oracle bugs, and 3 Microsoft bugs...
>
> The day began with Richard Zhu (fluorescence) targeting Apple Safari with a sandbox escape. Unfortunately, he could not get his exploit chain working within the time allotted due to a failure in the heapspray technique. Despite this, the bugs he brought to the contest were certainly interesting and were purchased through the regular ZDI program.
>
> Undaunted, Richard later returned to target Microsoft Edge with a Windows kernel elevation of privilege (EoP) exploit, bringing a flair for the dramatic: After his first attempt failed, he proceeded to debug his exploit in front of the crowd while still on the clock. His second attempt nearly succeeded... but the target blue screened just as his shell started. Oh, Richard! But his third attempt succeeded with one minute and 37 seconds left on the clock. In the end, he used

two use-after-free (UAF) bugs in the browser linked to an integer overflow in the kernel to successfully run his code with elevated privileges. Nice going, Richard! His dramatic last-ditch effort earned him $70,000 and 7 points towards Master of Pwn.

Next up, Niklas Baumstark (_niklasb) from the phoenhex team targeted Oracle VirtualBox. Apparently not one for added intrigue, his exploit immediately popped not one, but three!, different calculator apps to rub his success into Oracle's face. His demonstration qualified as a partial success as he used an Out-of-Bounds (OOB) read and a Time of Check / Time of Use (toctou) to still earn him $27,000 and 3 Master of Pwn points. It was a great demonstration, and we look forward to more of Niklas' research in the future.

The final attempt on Day One saw Samuel Groß (5aelo) of phoenhex targeting Apple Safari with a macOS kernel elevation of privilege (EoP). Last year, his exploit involved a touchbar component, and this year proved to be no different. Sam deftly leveraged a combination of a Just-In-Time optimization bug in the browser (Whoops, Apple!), then a macOS logic bug to escape the sandbox, and finally a kernel overwrite to execute code with a kernel extension... to successfully exploit Apple Safari. This lovely exploit chain earned him $65,000 and 6 points towards Master of Pwn. And as he did last year, he left a message for us on the touchbar once he was complete.

Day One also saw multiple entrants withdraw from the contest at the last minute -- some due to the multiple security patches released on Patch Tuesday, and others simply didn’t finish their exploit chains in time for the contest. We’re still in touch with these researchers and hope to acquire these bugs through regular ZDI purchase methods.

Day Two:

The second and final day of Pwn2Own 2018 concluded with an additional $105,000 USD and 11 more Master of Pwn points awarded...

The day began with the return of Richard Zhu (fluorescence), this time taking aim at Mozilla Firefox with a Windows kernel elevation of privilege (EoP). He eschewed all drama today and successfully popped Mozilla Firefox on his first attempt with his clever use of an out-of-bounds (OOB) write in the browser followed by an integer overflow in the Windows kernel to earn himself another $50,000 and 5 more Master of Pwn points... to bring his event total to $120,000 with a commanding lead for Master of Pwn.

Next up, Markus Gaasedelen (gaasedelen), Nick Burnett (itszn13), and Patrick Biernat of Ret2 Systems, Inc. targeted Apple Safari with a macOS kernel elevation of privilege (EoP). After experiencing some unexpected failures, they successfully demonstrated their exploit on the fourth attempt. Unfortunately, the contest rules only allow three attempts, so this counted as a technical failure. Still, the bugs used were purchased and disclosed to the vendor through the normal ZDI process while not counting toward the contest winnings.

The final entry of the day and of the contest saw a team from MWR labs – Alex Plaskett (AlaxJPlaskett), Georgi Geshev (munmap), and Fabi Beterke (pwnfl4k3s) - target Apple Safari with a slick and gritty sandbox escape. They snuck a heap buffer underflow past Safari when it

wasn't looking, then followed that up with an uninitialized stack variable in macOS to escape Safari's sandbox and gain code execution. In doing so, they earned $55,000 and 5 Master of Pwn points. Nice work, fellas!

And with that, this years' Pwn2Own came to a somewhat lackluster end. Richard Zhu was crowned Master of Pwn which earned him $120,000 over the two days as he accumulated 12 Master of Pwn points. Richard has participated in previous Pwn2Own contests, and we certainly hope he returns in the future to defend his title.

In total $267,000 was awarded during the two-day contest while revealing five Apple bugs, four Microsoft bugs, two Oracle bugs, and one Mozilla bug. Though smaller than some of our previous competitions, the quality of research was still extraordinary and served to highlight the difficulty of producing fully-functioning exploit for modern browsers and systems. We want to congratulate all those who participated in this year’s event. We also want to thank the multiple people who registered for the contest but needed to withdraw.

Meanwhile, DEF CON China will be May 11-13, 2018 in Beijing, China
https://www.defcon.org/html/defcon-china/dc-cn-index.html

So, overall, the security of Edge, Firefox, Safari, and VirtualBox were breached, and those vulnerabilities will be fixed.

One has to wonder, though, given prior years' much more dramatic results, what additional vulnerabilities the Chinese hackers have found and would have revealed had their government permitted their participation?

~30~